



**AFRL-OSR-VA-TR-2013-0539**

**(MURI-08) MODELING SYNERGIES IN LARGE HUMAN-MACHINE  
NETWORKED SYSTEMS**

**KATIA SYCARA**

**CARNEGIE-MELLON UNIVERSITY**

**09/25/2013**

**Final Report**

**DISTRIBUTION A: Distribution approved for public release.**

**AIR FORCE RESEARCH LABORATORY  
AF OFFICE OF SCIENTIFIC RESEARCH (AFOSR)/RSL  
ARLINGTON, VIRGINIA 22203  
AIR FORCE MATERIEL COMMAND**

<b>REPORT DOCUMENTATION PAGE</b>			<i>Form Approved</i> <b>OMB No. 0704-0188</b>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. <b>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</b>				
<b>1. REPORT DATE (DD-MM-YYYY)</b> 23-09-2013		<b>2. REPORT TYPE</b> Final Performance Report		<b>3. DATES COVERED (From - To)</b> July 2008-June 2013
<b>4. TITLE AND SUBTITLE</b> Modeling Synergies in Large Human-Machine Networked Systems			<b>5a. CONTRACT NUMBER</b> FA95500810356	
			<b>5b. GRANT NUMBER</b>	
			<b>5c. PROGRAM ELEMENT NUMBER</b>	
<b>6. AUTHOR(S)</b> Sycara, Katia, (PI); Lebiere, Christian;; Lewis, Michael; Cummings, Mary; How, Jonathan; Campbell Mark; Scerri, Paul; Parasuraman, Raja			<b>5d. PROJECT NUMBER</b>	
			<b>5e. TASK NUMBER</b>	
			<b>5f. WORK UNIT NUMBER</b>	
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> Carnegie Mellon University, 5000 Forbes Av. Pittsburgh, PA 15213 University of Pittsburgh, 4200 Fifth Av. Pittsburgh PA 15213 MIT ,77 Massachusetts Av. Cambridge MA 02139. Cornell University: Ithaca, NY 14850 George Mason University, 4400 University Dr. Fairfax, VA 22030			<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>	
<b>9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> Air Force Office of Scientific Research 875 North Randolph Street Arlington, VA. 22203			<b>10. SPONSOR/MONITOR'S ACRONYM(S)</b>	
			<b>11. SPONSOR/MONITOR'S REPORT NUMBER(S)</b>	
<b>12. DISTRIBUTION / AVAILABILITY STATEMENT</b> Approved for Public Release				
<b>13. SUPPLEMENTARY NOTES</b>				
<b>14. ABSTRACT</b> Network centric military systems (NCW) involve hundreds to thousands of manned and autonomous entities cooperating to achieve complex joint objectives in incomplete information environments. The overall goal of this multidisciplinary research is to provide validated theories and models, grounded in experiments with human operators that allow descriptive and predictive characterization of important properties and performance of complex and large-scale human-machine networked systems. The most significant results of the research were: (a) a scalable cognitive model framework that provides scalability while maintaining targeted cognitive fidelity (b) algorithms for automated large scale path planning robot systems, (c) predicting behavior, including vulnerabilities, of large scale heterogeneous complex networks, (d) algorithms for constrained multi-robot task assignment (e) scalable models of human robot control for independently operating robots, (f) robot self-reflection and queuing algorithms to schedule operator attention, (g) scalable displays, (h) models of human-robot decision making,, (j) models for planning and resource allocation in multi-robot teams with formal performance guarantees, and (k) human-automation collaborative scheduling				
<b>15. SUBJECT TERMS</b> complex networks of humans and machines, human supervisory control of multi-robot systems, scalable cognitive models, sched				
<b>16. SECURITY CLASSIFICATION OF:</b>			<b>17. LIMITATION OF ABSTRACT</b>  UU	<b>18. NUMBER OF PAGES</b>  92
<b>a. REPORT</b> U	<b>b. ABSTRACT</b> U	<b>c. THIS PAGE</b> U		
				<b>19b. TELEPHONE NUMBER (include area code)</b> 412-268-8825

# FINAL PERFORMANCE REPORT

**PI:** Katia Sycara

**Organization:** Carnegie Mellon University, 5000 Forbes Av. Pittsburgh, PA 15213

**Award Number:** Award Number: FA95500810356

**Award Title:** Modeling Synergies in Large Human-Machine Networked Systems

**Period of Performance:** 7/1/2008-6/30/2013

**Abstract:** Network centric military systems (NCW) involve hundreds to thousands of manned and autonomous entities cooperating to achieve complex joint objectives in incomplete information environments. The overall goal of this multidisciplinary research is to provide validated theories and models, grounded in experiments with human operators that allow descriptive and predictive characterization of important properties and performance of complex and large-scale human-machine networked systems. The most significant results of the research were: (a) a scalable cognitive model framework, ACT-UP, an effective abstraction of the ACT-R cognitive modeling system that provides scalability while maintaining targeted cognitive fidelity to aspects relevant to the application, (b) algorithms for automated path planning of large scale (hundreds) robot systems, (c) understanding and predicting behavior, including potential vulnerabilities, of large scale heterogeneous complex networks, (d) algorithms for constrained multi-robot task assignment (e) models of human performance as number of robots scale for independently operating robots, (f) robot self-reflection and novel queuing algorithms for scheduling operator attention, (g) scalable displays, (h) models of human-robot decision making, (i) models of human team interaction with automation, (j) models for planning and resource allocation in multi-robot teams with formal performance guarantees, and (k) human-automation collaborative scheduling.

## Summary of the Significant Work Accomplished

### 1. Scalable Cognitive Models (Lead: CMU-Psychology)

#### 1.1 Introduction

The ubiquitous and complex nature of information networks comprised of human and machine agents makes it essential to develop a methodology for their study that integrates the principles of behavioral research with the scalability of computational simulations. Therefore it is important to develop scalable cognitive models to allow studies in characteristics and performance of man-machine networked systems. This is significant since the availability of a scalable, easy-to-integrate, cognitively validated agent framework would make cognitive techniques accessible to a much broader range of potential users and applications.

The performance of teams is vital to the function of organizations. For instance, small and large teams of warfighters may be united in pursuing overall goals and trained to precisely interact with their environment according to defined protocols. Yet, achieving an information advantage is crucial. Do they exchange vital information expediently and reliably? How is such communication organized? How are joint decisions taken? Such questions have been investigated using simple if not simplistic computational games and multi-agent simulations. Recent advances in cognitive modeling provide a high-fidelity account of individual performance. Recent breakthroughs establishing the science of networks allow us to describe the structural properties of teams, and propose mechanisms that may lead to the creation of team structures as

we observe them. The ubiquitous and complex nature of information networks makes it essential to develop a methodology for their study that integrates the principles of behavioral research with the scalability of computational simulations. Our interdisciplinary approach combining multi-agent simulation, cognitive modeling and network science yields new insights in the function of teams and the emergence of communication systems.

CMU Psychology has led the investigation of scalability in robust cognitive models in order to explain and predict team behavior and the emergence of joint communication and action. Through a new implementation of the ACT-R theory (Anderson 2007), called ACT-UP, we mitigate the tradeoff between fidelity and complexity in cognitive modeling, providing a faster, rapid-prototyping environment. This has been applied in a series of cognitive models in the teamwork domain. Furthermore, we have turned to empirical validation of these models

In one application of these new methods, we conduct the first large-scale experiments with synchronous team interaction in a controlled environment and with a well-defined, algorithmically analyzable, communication-dependent task. Prior work has investigated team interaction using simple games (Kearns lab, U Penn and Winter and Watts, Yahoo Research), but has avoided the use of communication or more complex task dynamics. Our team with an interdisciplinary background in Computer Science, Cognitive Psychology and Linguistics was suited to design these new simulations and experiments.

## **1.2 Robust and efficient large-scale cognitive modeling in the ACT-UP framework.**

Work on all models in this MURI has benefited from a novel common implementation of the ACT-R theory (Anderson 2007). *ACT-UP* is an abstraction of ACT-R designed to provide the following advantages:

- speed up development time by focusing programmer efforts
- scalability to large numbers of agents for network simulations
- targeted cognitive fidelity only to aspects relevant to the task
- facilitate integration with other programming/modeling frameworks

ACT-UP achieves those objectives by providing a direct API to the key aspects of ACT-R functionality, such as memory retrieval, production matching, visual search, etc. This API approach allows the modeler to leverage only the aspects of the architecture relevant to a given application, thus speeding up development time as well. The lightweight framework, as opposed to the commitment required by a monolithic architecture, provides scalability to large numbers of agents and easy integration with other programming or modeling languages

ACT-UP provides an opposite solution to another approach to providing a higher-level cognitive language, the High-Level Behavioral Representation Language (HLSR): while HLSR attempts to abstract away from the key architectural components, ACT-UP exposes them directly. But while HLSR still commits to running the full model within the architectural framework, ACT-UP only commits to running the key elements and allows the modeler to abstract the other ones for tractability or efficiency.

Experience with the Language Evolution model (see below) shows that ACT-UP can provide scalability and efficiency in two ways. *Simulation Scalability*: We observed a speed-up of an estimated 1,000% in the multi-agent simulation of language evolution compared to an earlier ACT-R implementation of the model, owed to the new implementation but also to the fact that underspecified aspects of the task model can be executed much more efficiently. *Modeling Scalability*: a modeling effort of about two person-months in ACT-R translated, in this case study,

to one person-week in ACT-UP, again due to better prototyping and debugging facilities, but also due to under-specification and shorter turnaround-times.

ACT-UP, a re-implementation of the ACT-R theory that introduces high-level, high-fidelity modeling, rapid prototyping, and better scalability, has been made available for the scientific community. The system has been validated thoroughly by re-implementing known ACT-R models and verifying the results.

David Reitter and Christian Lebiere. Accountable modeling in ACT-UP, a scalable, rapid-prototyping ACT-R implementation. In *Proceedings of the 10th International Conference on Cognitive Modeling (ICCM)*, pages 199-204, Philadelphia, PA, 2010.

Christian Lebiere, Andrea Stocco, David Reitter, and Ion Juvina. High-fidelity cognitive modeling to real-world applications. In *Proceedings of the NATO Workshop on Human Modeling for Military Application*, Amsterdam, NL, 2010. 19 pages.

### **1.3 The Geo Game Experimental Framework**

We designed and implemented an experimental framework, called the Geo Game series of experiments that study collaboration and communication in networked human groups. The Geo Game is a foraging game, developed by CMU-Robotics, CMU-Psychology and the U Pittsburgh teams, that is designed to exercise individual cognitive abilities, specifically memory, perceptual and communication capacities. As such, the task exercises human abilities typically required in real-life teamwork tasks, as well as team-specific skills.

In the Geo Game, participants have to locate hidden items scattered throughout a virtual world represented by a map. Exploring the map is time-consuming, but they may communicate their findings using written messages, greatly speeding up their work. Individuals only communicate with a predefined subset of teammates. In the current series of experiments, we use small-world networks to define the communication paths, whose structure is representative of larger human and non-human communication and cooperation networks.

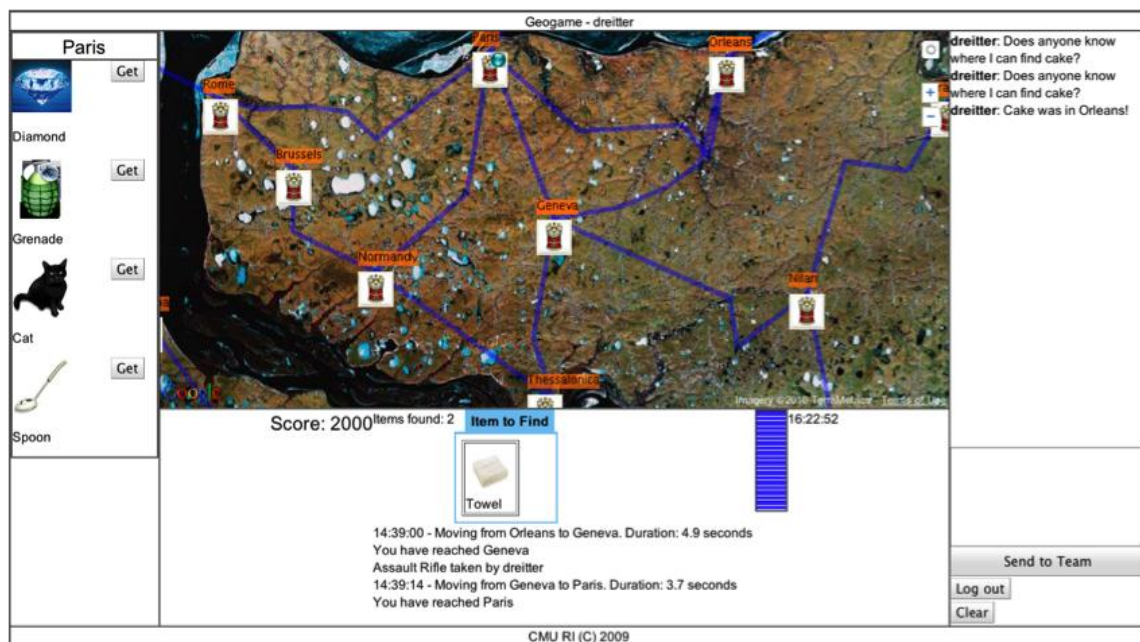
In any such real-world task, communication and task execution are usually co-dependent, yet represent a tradeoff: communication takes time and attentional resources from the main objective. We present a cognitive model of an experimental task consisting of a collaborative and competitive game played by groups of human participants organized in a small-world graph.

Through a range of possible manipulations, the Geo Game platform allows us to answer questions about how information propagates in networks, how it modulates the interaction of adversarial networks, how it is acquired and retained by networks (accommodating individual limitations), how communication mechanisms are developed and optimized by communities, and how controlling some of these parameters through technical means can improve task success.

Our initial experiments investigate techniques to optimize collaboration. The first experiments with the Geo Game concerned communication policies for individuals working in teams. The experiment involved teams (20 participants per team) of humans playing a cooperative game. The effect of local communication policies on the efficiency and the performance of networked participants was observed. The model follows the ACT-R theory and provides a formalization of the decision-making processes and tradeoffs involved.

Specifically, we looked at the use of *communication policies* in networks, hypothesizing that judicious communications not only are more effective overall, but also more efficient. The initial experiments confirmed this, and they also provided suggestive evidence that individuals that communicate with only a few others in the network benefit more from a policy of judicious, targeted communications than do the well-connected ones.

In further experiments, we found results consistent with substantial adaptivity among subjects. Some subject groups were able to perform well even under the non-targeted, “information overload” condition; in a control condition, we were able to obtain good performance also from subjects who did not communicate at all. Post-experiment interviews suggested that subjects were able to memorize information that helped them play the game.



**Figure 1: The Geo Game screen showing the participants’ task graphical interface. In particular, the figure shows the map with city names, a panel in the left hand side showing articles that are currently in Paris (to get this the participant has to “go to” Paris), a chat interface below the map, a window that shows the item that the participant has to find (Towel), a panel in the left hand side showing requests and replies from various team members of the participant.**

Once communication takes place, information state is maintained by individuals, but also non-redundantly by the network. Much of our work related to information state maintenance in human individuals and human or human-machine networks. In one of these studies, a multi-model simulation, two *information decay* methods were examined that help multi-agent systems cope with dynamic environments. The agents in this simulation have human-like memory and a mechanism to moderate their communications: they forget internally stored information via temporal decay, and they forget distributed information by filtering it as it passes through a communication network. The agents play a foraging game, in which performance depends on communicating facts and requests and on storing facts in internal memory. Parameters of the game and agent models are tuned to human data. Agent groups with moderated communication in small-world networks achieve optimal performance for typical human memory decay values, while non-adaptive agents benefit from stronger memory decay. The decay and filtering strategies interact with the properties of the network graph in ways suggestive of an evolutionary co-optimization between the human cognitive system and an external social structure.

David Reitter and Christian Lebiere. Towards cognitive models of communication and group intelligence. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*, pages 734-739, Boston, MA, July 2011.

David Reitter, Katia Sycara, Christian Lebiere, Yury Vinokurov, Antonio Juarez, and Michael Lewis. How teams benefit from communication policies: information flow in human peer-to-peer networks. In *Proceedings of the 20th Behavior Representation in Modeling & Simulation (BRIMS)*, 2011.

#### **1.4.Cognitive models of distributed network interaction**

Using the ACT-UP cognitive modeling toolkit, we have developed cognitive models for a number of specific network activities, including spatial path planning and navigation in multi-robot control systems, language evolution, and control and decision-making. Finally, we developed an integrated model of these cognitive activities in the context of the Geo Game foraging simulation to validate their interaction in the context of a complex task.

##### **1.4.1. Spatial path planning in mazes, multi-robot control systems, and general navigation tasks**

Planning a path to a destination, given a number of options and obstacles, is a common task. We developed a two-component cognitive model that combines retrieval of knowledge about the environment with search guided by visual perception. In the first component, subsymbolic information, acquired during navigation, aids in the retrieval of declarative information representing possible paths to take. In the second component, visual information directs the search, which in turn creates knowledge for the first component. The model is implemented using the ACT-UP cognitive toolkit and makes

realistic assumptions about memory access and shifts in visual attention. We derived simulation results for memory-based high-level navigation in grid and tree structures, and visual navigation in mazes, varying relevant cognitive (retrieval noise, visual finsts) and environmental (maze and path size) parameters.

We applied and evaluated the model in an experiment involving visual path planning for multiple, remote robots in a partially visible building, with a partial 2D map available. Participants in the experiment defined waypoints for each robot to circumnavigate obstacles and explore the building. Our visual planning model is evaluated using the experimental data with a normalized metric of the fit between model and subject itineraries. Through model fit we observe individual differences in strategies to cope with task demands.

David Reitter and Christian Lebiere. A subsymbolic and visual model of spatial path planning. In: *Proc. Behavior Representation in Modeling and Simulation (BRIMS)*, 2009. Best paper award BRIMS 2009.

David Reitter, Christian Lebiere, Michael Lewis, Huadong Wang, and Zheng Ma. A cognitive model of visual path planning in a multi-robot control system. In: *Proceedings Systems Man Cybernetics 2009 (IEEE-SMC)*, San Antonio, TX, 2009.

David Reitter and Christian Lebiere. A cognitive model of spatial path planning. *Computational and Mathematical Organization Theory*, 16(3):220-245, 2010.

#### **1.4.2. Towards explaining the evolution of domain languages with cognitive simulation**

We simulated the evolution of a domain language in small speaker communities. Data from published experiments show that human communicators can evolve graphical languages quickly in a constrained task (Pictionary), and that communities converge towards a common language even in the absence of feedback about the success of each communication. We postulated that simulations of such horizontal evolution have to take into account properties of human memory (cue-based retrieval, learning, decay). We implemented a model that can draw abstract concepts through sets of non-abstract, related concepts, and recognize such drawings. The knowledge base is a network with association strengths randomly sampled from a natural distribution found in a text corpus; it is a mixture of knowledge shared between agents and individual knowledge. In three experiments, we showed that the agent communities converge, but that initial convergence is stronger when communities are structured so that the same pairs of agents interact throughout. Convergence is weaker in communities when agents do not swap roles (between recognizing and drawing), predicting the necessity of bi-directional communication in domain language evolution. Average and ultimate recognition performance depends on how much of the knowledge agents share initially.

Originally, we developed this model according to previously available data for small (8-person) communities. In following years, this model has been integrated in a network-based simulation with up to 1,000 cognitive models, which interact to develop a common vocabulary. Contrasting a range of networks that differed by their structural form, we found striking differences between organizational hierarchies (trees) and naturally occurring small world networks. While trees performed the task best due to excellent



local convergence, they greatly suffered when the network was reconfigured. In other words, they did not show global convergence, and such teams failed to develop a common language. Instead, they developed many “local” languages. Small worlds performed well in the task and maintained their performance across configurational changes. Thus, small world networks represent a more robust form of organization with respect to tasks that depend on the exchange of information via language.

We developed a cognitive model of an experimental task consisting of a collaborative and competitive game played by groups of human participants organized in a small-world graph. In an experiment involving teams of humans playing a cooperative game, the effect of local communication policies on the efficiency and the performance of networked participants was observed. A simulation of the hypothetical case of unnatural memory decay shows decreased performance and supports a prediction of the thesis that memory limitations have co-evolved with social structure. In a more advanced line of work, we cast decay in individual memory to explain a complex pattern of linguistic adaptation effects that explain how small or large teams of people effortlessly align their languages. The psycholinguistic literature has identified two such syntactic adaptation effects in language production: rapidly decaying short-term priming and long-lasting adaptation. To explain both effects, we developed a model of syntactic priming that applies a wide-coverage linguistic theory that explains priming as a standard memory effect. In this model, two well-established mechanisms, base-level learning and spreading activation, account for long-term adaptation and short-term priming, respectively. Our model simulates incremental language production and in a series of modeling studies we show that it accounts for a pattern of empirically documented results. An understanding of the cognitive mechanisms of adaptation in language use are relevant for the development of human-computer interfaces, for communication protocols within teams of humans and mixed human-machine teams.

David Reitter and Christian Lebiere. Towards explaining the evolution of domain languages with cognitive simulation. In: *Proceedings of the 9th International Conference on Cognitive Modeling (ICCM)*, Manchester, UK, 2009.

David Reitter and Christian Lebiere. Did social networks shape language evolution? A multi-agent cognitive simulation. In *Proc. Cognitive Modeling and Computational Linguistics Workshop (CMCL)*, pages 9-17, Uppsala, Sweden, 2010. Association for Computational Linguistics.

David Reitter and Christian Lebiere. On the influence of network structure on language evolution. In Ron Sun, editor, *Proc. CogSci Workshop on Cognitive Social Sciences: Grounding the Social Sciences in the Cognitive Sciences*, Portland, Oregon, 2010.

David Reitter and Christian Lebiere. How groups develop a specialized domain vocabulary: A cognitive multi-agent model. *Cognitive Systems Research*, 12(2):175-185, 2011.

David Reitter, Frank Keller, and Johanna D. Moore. A computational cognitive model of syntactic priming. *Cognitive Science*, 35(4), p.587-637. 2011.

David Reitter. Lexical language evolution in networked human groups. In *Words and Networks: Language Use in Socio-Technical Networks (WON 2012)*, Chicago, IL, 2012.

### 1.4.3. A two-level, multi-strategy model of memory-based control

Multi-tasking, high demand environments often require human operators to balance dynamic, strategic coordination tasks (including communication) with real-time control demands (such as driving a vehicle or flying an aircraft). We developed a model of real-time control that can determine the external and internal factors that affect human performance, such as input-response feedback delays (external) or altered memory performance (internal). Real-time control is a common task to humans, whose performance improves with experience. Control tasks are usually similar in their general structure. ). We developed the model in the context of the Dynamic Stocks and Flows (DSF) cognitive modeling competition to take advantage of human data available for a number of conditions and test the predictiveness of the model in unseen conditions. In the DSF task, human subjects iteratively control water flow out of a water tank, reacting to a changing, independently determined inflow to the tank. Thus, the core task is to estimate the development of the inflow from discrete samples; the distribution underlying the inflow is a function of time or iterative steps. Once the next inflow is estimated, subjects can counteract it by choosing an appropriate outflow valve setting. (This corresponds to the real-world task of maintaining altitude and airspeed when piloting an aircraft subject to external factors.) In the empirical data available to design the model, the inflow function was manipulated across four conditions, combining linear and non-linear, decreasing and increasing inflow. Our model attempts to bridge the specifics of the experiment that produced the provided data, which involved a learning process and arithmetic decision-making, and real-life control problems, which also involve less discrete, non-arithmetic strategies to react to incremental environmental changes and to correlations of human actions and delayed environmental effects. Our proposed control model thus had two layers: a meta-cognitive level, choosing an optimal strategy to address the problem, and a task-specific level, which executes each strategy. The model won the DSF competition by providing the best generalization to undisclosed experimental manipulations, such as fluctuating inputs and outputs characteristic of an unstable environment, and control delays reflecting the complexity of the underlying system.

We also applied a similar modeling approach to another agent modeling competition, the Lemonade Game. The Lemonade Game is a three-player game in which players have to pick locations on a circular board, which are as far away as possible from those chosen independently by other players. Players may observe other player's moves and infer their strategies. The game was studied using a competition of cognitively motivated agents, which inherit properties of adaptivity and stochasticity from human memory and decision-making, and simplistic, yet effective, agents implementing fixed strategies. Our model demonstrated that metacognition constitutes the unique attribute that allows sophisticated agents to adapt to unforeseen conditions, cooperators and competitors.

David Reitter, Ion Juvina, Andrea Stocco, and Christian Lebiere. Resistance is futile: Winning lemonade market share through metacognitive reasoning in a three-agent cooperative game. In *Proceedings of the 19th Behavior Representation in Modeling & Simulation (BRIMS)*, Charleston, SC, 2010.

- Kevin A. Gluck, Clayton T. Stanley, Jr. L. Richard Moore, David Reitter, and Marc Halbrügge. Exploration for understanding in model comparisons. *Journal of Artificial General Intelligence*, 2(2):88-107, 2010.
- David Reitter. Metacognition and multiple strategies in a cognitive model of online control. *Journal of Artificial General Intelligence*, 2(2):20-37, 2010. Winning entry of Dynamic Stocks and Flows cognitive modeling competition.
- Christian Lebiere and John R. Anderson. Cognitive constraints on decision making under uncertainty. *Frontiers in Cognition* 2 (305). 2011.

#### **1.4.4. Information foraging in the Geo Game simulation**

To model human performance in the Geo Game experimental framework, we have developed a scalable, cognitively valid agent simulation comprising ACT-UP and a network library that makes cognitive techniques accessible to a broader range of potential users and applications, and is currently seeing re-use in our own groups. A cognitive simulation of the Geo Game was implemented using the ACT-UP system. In this simulation, a number of instances of a cognitive model play the Geo Game; the simulation obtains task performance similar to that of human performance. This simulation not only explains some of the results obtained experimentally, but it also allows us to predict the effects of further manipulations. For instance, we used the simulation to decide which aspects of the game to control and keep constant across experimental groups and conditions, and which aspects to randomize. This question is highly relevant in complex, dynamic experiments like ours. We are unaware of previous work predicting the effect of randomization in multi-subject experiments with dynamic tasks.

We hypothesize that individual cognition has co-evolved with social structure to allow the individual to externalize memory in a robust storage mechanism, to optimize the development of a common communication system (e.g., vocabulary) and ultimately to perform well. Large-scale cognitive modeling allowed us to test that hypothesis. Concretely, simulations that manipulate architectural parameters have shown that typical values for memory performance that have been empirically validated in the ACT-R literature also result in good performance in the Geo Game model. The key research issue involved is the fundamental tradeoff between the costs and benefits of information acquisition and processing. The basic assumption of the development of information and communication infrastructure is that more information is better. Our research approach is two-pronged: experimentally investigate the impact of that tradeoff on performance, and model the cognitive and perceptual processes by which it takes place, including attentional and adaptive mechanisms. The goal is to develop an understanding that allows the design of systems that achieve the best possible performance given technical and cognitive limitations.

The Geo Game provides a unique platform for experimentation of information rich tasks in networked situations. ACT-UP is a modeling toolkit that allows for the lightweight, scalable integration of human performance models in networked simulations. Together, they provide an approach to modeling and simulation that can be used to evaluate and design a broad range of information systems in networked settings. In any such real-

world task, communication and task execution are usually co-dependent, yet represent a tradeoff: communication takes time and attentional resources from the main objective. Once communication takes place, information state is maintained by individuals, but also non-redundantly by the network. Much of our work related to information state maintenance in human individuals and human or human-machine networks. In one of these studies, a multi-model simulation, two information decay methods were examined that help multi-agent systems cope with dynamic environments. The agents in this simulation have human-like memory and a mechanism to moderate their communications: they forget internally stored information via temporal decay, and they forget distributed information by filtering it as it passes through a communication network. The agents play a foraging game, in which performance depends on communicating facts and requests and on storing facts in internal memory. Parameters of the game and agent models are tuned to human data. Agent groups with moderated communication in small-world networks achieve optimal performance for typical human memory decay values, while non-adaptive agents benefit from stronger memory decay. The decay and filtering strategies interact with the properties of the network graph in ways suggestive of an evolutionary co-optimization between the human cognitive system and an external social structure.

- David Reitter and Christian Lebiere. Towards cognitive models of communication and group intelligence. In *Proc. 33rd annual meeting of the Cognitive Science Society*, Boston, MA, 2011.
- David Reitter and Christian Lebiere. Social cognition: Memory decay and adaptive information filtering for robust information maintenance. In *Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12)*, 2012.
- David Reitter and Paul Scerri. Social multi-agent learning with simple and cognitive agents. In *Proceedings of CAOSS 2012: Workshop on Computational and Online Social Science*, New York, N.Y., 2012.
- David Reitter and Paul Scerri. Smooth dynamics, good performance in cognitive-agent congestion problems. In *Proceedings of the 35th Annual Meeting of the Cognitive Science Society*, 2013.
- Paul Scerri and David Reitter. Cognitive instance-based learning agents in a multi-agent congestion game. In *Workshop on Information Sharing in Large Scale Multi-Agent Systems, at AAMAS 2013*, 2013.

## **2. Large Scale Multi Robot Path Planning Algorithms (Lead: CMU-Robotics)**

In many domains, teams of hundreds of agents must coordinate together to plan on performing tasks in a complex environment. Naively, this could require that agents take in every teammate's states, observations, and choice of actions into account when making decisions about their own actions. This results in a huge joint space over which it is computationally intractable to find solutions. In certain problems, however, searching this complete space may not be necessary. We have studied methods to substantially reduce the search space of joint planning problems for teams of agents in domains where individual agents often act independently, but there are certain combinations of states and actions where two or more agents share a non-factorable transition, reward, or observation functions. Previous work has exploited knowledge of this type of structure to reduce the search space of a centralized joint policy search. However, in our work, teams

are assumed to be very large, consisting of hundreds of agents, and thus additional techniques to reduce search complexity are needed.

In order to handle such large team sizes, we exploit two particular properties of our domains of interest. First, although there can be a large number of interactions that are possible, it is often the case in these domains that the number of interactions that actually occur in any given solution instance is quite small. By dynamically discovering relevant interactions rather than trying to handle every possibility, algorithm convergence can be greatly improved. Second, in many domains, computational power itself is distributed across a team of agents. This means that within these domains, running a planner requires either that the algorithm is inexpensive enough to run on a single agent, or fully distributable. Thus, we focus on distributed approaches that have access to computational resources that grow linearly with team size in comparison to centralized approaches, making them much easier to scale to very large teams.

We have addressed two planning problems in which these characteristics occur: multiagent path planning and Distributed POMDPs with Coordination Locales (DPCLs), a subproblem of the canonical Dec-POMDP. Using similar approaches of dynamically detecting and resolving interactions, we are able to adapt existing solution techniques to significantly improve scalability.


In the former case, with agents planning simultaneous paths over a grid structure, the result is Distributed Prioritized Planning (DPP), a simple variant of the sequential Prioritized Planning. Results with DPP demonstrate that iterative planning in situations where interaction is sparse can produce efficient solutions in relatively little iteration with respect to team size. However, they also emphasize the importance of low variance in individual agent planning times in allowing distributed, iterative planning to be more effective than sequential, decoupled planning.

The latter problem, DPCL, is addressed by the more powerful D-TREMOR algorithm, an extension to the centralized, iterative TREMOR algorithm. D-TREMOR significantly scales the TREMOR algorithm by replacing joint search and evaluation steps with fully distributed heuristic approximations. Performance is demonstrated in solutions of DPCLs with over 100 agents in a simplified rescue domain. The results show the efficacy of prioritization and randomization in adjusting models of teammates' actions for the interactions modeled in the rescue domain, but suggests that additional work is necessary to further improve performance and generalize D-TREMOR to other potential types of agent interactions.

D-Tremor provides a tractable model for multiagent sequential decision making problems. By constraining interactions between agents to have symmetric and idempotent effects, and specifically defining those effects for each agent, our distributed POMDP algorithm (called RDPCL) is easily specifiable for many agents while remaining computationally tractable. We implemented different instantiations of this approach and compared performance. The algorithm is capable of planning policies for more than 100 agents. The ability to represent interesting problems has been made dramatically more powerful and the heuristics to get good algorithm convergence have been developed. The



**Figure 2: A sample map solved by the DPP algorithm. Agents start at each of the circles on the map, and must reach their matching star positions without colliding with any other agents.**



**Figure 3: A sample map solved by the D-TREMOR algorithm. Rescue and cleaner robots start at the marked locations and must coordinate to rescue as many victims as possible**

algorithms have been tested in multiple domains and are now being transitioned to a HSBC program for contingency planning in complex, multi-actor environments.

We identified dynamic sparsity as a characteristic of many multiagent decision problems. Dynamic sparsity is a powerful structural property in many planning problems, greatly restricting the joint interactions between agents given their policies. Our DIMS framework exploits dynamic sparsity directly, using iterative solving to restrict necessary policy computation to only interactions which arise during the planning process rather than all the possible interactions.

We developed *model shaping* heuristics for distributed planning. In order to reach good solutions, we introduced priority and randomization heuristics to quickly and reasonably resolve interactions. We demonstrated that by adding randomization to our DIMS framework improves solution quality at the expense of determinism, while adding prioritization improves determinism at the expense of optimality.

In collaboration with the University of Pittsburgh, we defined two benchmark problems, the rescue domain and the convoy domain that mirror real world applications. These problems are well defined for any number of agents and contain complex agent interactions both negative interactions (eg collisions) and positive interactions (e.g one robot fulfilling preconditions for another one to act).

P. Velagapudi, K. Sycara, and P. Scerri, Decentralized prioritized planning in large multirobot teams, In IROS'10, 2010.

P. Velagapudi, P. Varakantham, K. Sycara, and P. Scerri Distributed Model Shaping for Scaling to Decentralized POMDPs with Hundreds of Agents, In AAMAS'11, 2011.

Varakantham, P., Yeoh, W., Velagapudi, P., Sycara, K., Scerri, P. "Prioritized Shaping of Models for Solving DEC-POMDPs" International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-12), Valencia, Spain, June 4-8, 2012

### 3. Complex Networked Systems (Lead: CMU-Robotics)

#### 3.1. Emergent Information Dynamics

In the near future, large heterogeneous teams of robots, agents, and people will be utilized to solve problems in a variety of applications including search and rescue and the military. The sheer size of such teams will mean that the amount of data collected by the team will be overwhelming for its constituents. For this reason, team members will need to share concise information abstractions, i.e. conclusions, to maintain shared situational awareness.




Figure 4. Information cascade distribution  $P(c)$  where  $c$  is the cascade size.

The physics of communication, along with environmental constraints, will require team members to communicate via a point to point associates network. This will in turn lead to complex information dynamics and emergent phenomena, which in turn leads to unpredictability. Large heterogeneous teams will often be in situations where sensor data that is uncertain and conflicting is shared across a peer to peer network. Not every team members will have direct access to sensors. Thus team members will be influenced mostly by information of team mates with whom they communicate directly. We investigated the dynamics and emergent behavior of a large team sharing beliefs to reach conclusions about the world. We found empirically that the dynamics of information propagation in such belief sharing systems are characterized by *information cascades* of belief changes caused by a single additional sensor reading input to the system. The distribution of the size of these cascades dictates the speed and accuracy with which the team reaches conclusions. A key property of the system is that it exhibits qualitatively

different dynamics and system performance over small range of changes in the parameters of the system. In one particular range the system exhibits behavior known as *scale invariant dynamics* which we empirically find to correspond to dramatically more accurate conclusions being reached by the overall system. Due to the fact that the ranges are very sensitive to system configuration details, the parameter ranges over which specific system dynamics occur are extremely difficult to predict precisely. We have developed (a) techniques to mathematically characterize the dynamics of the team belief propagation, (b) obtain the relation between they dynamics and overall system performance and (c) developed a novel distributed algorithm that the agents in the team use locally to steer the whole system to areas of optimized performance. In particular, the




Figure 5: the x axis denotes “trust” in a neighbor (conditional probability of believing a neighbor) and the y-axis denotes reliability of agents’ conclusions. We see that at  $cp=0.67$  there is a dramatic performance peak.




Figure 6. By using the Theory of Branching processes, we developed an algorithm, DACOR that each agent uses locally to adaptively change the network dynamics to maintain high performance quality.



agents make local adjustments to conditional probabilities on neighbors observations that move the team towards the parameter ranges where scale invariant dynamics occur for any network type, thus dramatically improving system performance. This algorithm also minimizes disruption to the overall network making it practically applicable in real world systems. Our study shows that small amounts of anomalous information introduced to such a belief sharing system can cause errors on a system-wide scale due to the intrinsic dynamics of the system. This could potentially be exploited by a malicious agent attempting to disrupt such a system. Both analytical and empirical evidence is provided to support this assertion. Previous attempts to describe the vulnerabilities of complex networked system primarily focus on finding vulnerabilities in the network topology without consideration of the dynamics of the process taking place on the network. In our work, the dynamics on the network have a dramatic impact on the vulnerability of the system.. We showed that a team of agents could tune their local trust such that the frequency distribution of cascades of changes in belief followed a power law. When the team was tuned like this, the team's ability to rapidly reach correct conclusions despite noisy data and limited communications was shown to be dramatically higher. However, we show that when a system is tuned like that, it also becomes vulnerable to malicious attacks.

Glinton, R., Paruchuri, P., Scerri, P. Sycara, K “Self-organized criticality of belief propagation in large heterogeneous teams”, Hirsch, M Pardalos, P. and Murphy R (eds), Dynamics of Information Systems, Springer, 2009.

Glinton, R., Sycara, K, Scerri, P. “Exploiting Scale Invariant Dynamics for Efficient Information propagation in Large Teams”, Proceedings of the 2010 Conference on Autonomous Agents and Multi-Agent Systems, Toronto, CA, May, 2010 (*Second Place for Best Paper Award*).

Glinton, R., Scerri, P., Sycara, K. “An Explanation for the Efficiency of Scale Invariant Dynamics of Information Fusion in Large Teams”, International conference on Information Fusion (Fusion2010), July 26-29, Edinburgh, UK, 2010.

### **3.2. Vulnerabilities in Complex Networked Systems**

We conducted an analysis to show that for a system exhibiting scale invariant dynamics, a single anomalous sensor reading could result in a number of agents on the order of the size of the system coming to the incorrect conclusion. The analysis compares the rate at which the probability that an agent is on the edge of coming to a correct conclusion, called the percolation probability, increases relative to the same probability for an incorrect conclusion. The analysis reveals that these two numbers remain close until the agents in the system converge. Although this difference is biased towards correct conclusions, the analysis shows that this difference is small enough for a few anomalous sensor readings to push large numbers of agents towards incorrect conclusions. To confirm the predictions of the analysis we empirically explored the effect of injecting a single incorrect sensor reading into the system on the correctness of conclusions reached by agents in the system. We showed empirically by exhaustively searching trajectories of system execution that there is always a point in that trajectory where injecting a single sensor reading can lead to system wide incorrect conclusions. We further show that an

adversary could mount an effective attack on the system if the adversary had global knowledge of the distance of the system from the percolation threshold for the incorrect conclusion.

Just as complex systems can be attacked from external sources, it is also possible for attacks to originate from within. Thus it is necessary to understand the potential vulnerabilities of such a system to threats from within. To this end we studied the vulnerability of the agents within the system to reaching incorrect conclusions as a result of the action of Byzantine agents within the system. Specifically, we studied mechanisms for picking the most vulnerable points in the network for attack by Byzantine agents. We explored several different mechanisms for selecting which nodes are Byzantine, using methods typically employed in the study of the vulnerabilities in network topologies to network disintegration. The study reveals that the most effective method is that which selects the nodes with the maximum number of neighbors. Finally, our study shows that as the number of Byzantine agents in the network increases, the trust range between agents that results in a scale invariant distribution of cascades is no longer optimal. As the number of Byzantine agents increases the optimal value of trust is lowered slightly with the agents becoming slightly more conservative to account for the misinformation circulating in the system.

In a large distributed system it is unlikely that an adversary would have access to the global network state or topology, thus it is desirable to study whether an effective attack on the system could be launched using only local knowledge of the network state and topology. To investigate the feasibility of a practical attack we developed a local algorithm, where Byzantine agents use knowledge of the local connectivity and a local estimate of the percolation threshold to decide when and where to focus an attack. We found that such an attack is as effective, in reducing the number of agents that come to a correct conclusion, as an attack mounted with full knowledge of the system state and network topology.

Glinton, R., Scerri, P., Sycara K., An Investigation of the Vulnerabilities of Scale Invariant Dynamics in Large Teams Proceedings of the 2011 Conference on Autonomous Agents and Multi-Agent Systems, May 2-6, Taipei, Taiwan, 2011.

### **3.3. Multi-agent learning in large scale networked heterogeneous systems**

Building on previous work that showed the utility of scale invariant dynamics at reaching consensus, a multi-agent learning algorithm was developed with the same inspiration. By carefully modulating the rate at which agents communicate, the overall learning rate could be substantially improved, despite the non-stationary learning environment created by the simultaneous learning. In another strand of this work, our previous information sharing algorithms were extended to handle situations where the agents slowly changed from broadcast to peer-to-peer communication as they moved around the environment and needed to adjust their communication algorithms for best overall performance.

P. Scerri, "Modulating Communication to Improve Multi-Agent Learning Convergence", In OPTMAS'12 Workshop at AAMAS-12.

### **3.4 Non-Zero Sum Multiagent Network Security Games**

Moving assets through a transportation network is a crucial challenge in hostile environments such as future battlefields where malicious adversaries have strong incentives to attack vulnerable patrols and supply convoys. Intelligent agents must balance network costs with the harm that can be inflicted by adversaries who are in turn acting rationally to maximize harm while trading off against their own costs to attack. Furthermore, agents must choose their strategies even without full knowledge of their adversaries' capabilities, costs, or incentives. We modelled this problem as a non-zero sum game between two players, a sender who chooses flows through the network and an adversary who chooses attacks on the network. We advance the state of the art by: (1) moving beyond the zero-sum games previously considered to non-zero sum games where the adversary incurs attack costs that are not incorporated into the payoff of the sender; (2) introducing a refinement of the Stackelberg equilibrium that is more appropriate to network security games than previous solution concepts; and (3) using Bayesian games where the sender is uncertain of the capabilities, payoffs, and costs of the adversary. We provide polynomial time algorithms for finding equilibria in each of these cases. We also show how our approach can be used for games where there are multiple adversaries.

Okamoto, S., Hazon, N., Sycara, K. "Solving Non-Zero Sum Multiagent Network Flow Security Games with Attack Costs", International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Valencia, Spain, June 4-8, 2012.

Steven Okamoto\_, Praveen Paruchuri, Yonghong Wang, Katia Sycara, Janusz Marecki and Mudhakar Srivatsa "Multiagent Communication Security in Adversarial Settings", International Conference on Intelligent Agent Technology, Lyon, France, August 22-27, 2011.

## **4. Algorithms for Multi-Robot Task Assignment with Formal Guarantees (Lead CMU)**

In many multi-robot applications like environmental monitoring, search and rescue, disaster response, extraterrestrial exploration, the tasks that the robots need to perform are not known beforehand but arise as the robots are executing their missions. In such scenarios, robots may be able to do more than one task during a mission depending on their capabilities and battery life. Since battery life for a robot is limited there will be an upper bound on the number of tasks that a robot can do during a mission. The problem of allocating tasks to robots when the tasks are not known beforehand but may arise in an *online* fashion is called the *online task allocation (OTA) problem* or *online assignment problem*. Depending on the characteristics of the tasks and the capability of the robots, different versions of the OTA problem can be formulated. In the simplest version of OTA, also known as online maximum weight bipartite matching problem (MWBMP), the tasks arrive one at a time and each robot can do at most one task in the mission. Each robot-task pair has a certain payoff and the objective is to maximize the total payoff of

the multi-robot system. We study a generalization of the online MWBMP, where the tasks can arise dynamically in groups and each robot can do at most one task in each group, but can do more than one task in the whole mission. The abstract problem is motivated by two different kinds of scenarios arising in applications: (a) Tasks arise dynamically in groups, where each group consists of tightly-coupled tasks, i.e., tasks which robots must perform simultaneously, and thus each robot can only be assigned to one of them; (b) There exist group precedence constraints among tasks, i.e., only after the current group of tasks are all completed by robots, the subsequent group of tasks can get started, and the corresponding (payoff) information is revealed to robots. To fully explore the parallelism, each robot can be assigned to at most one task in each group to increase the efficiency. A special case, where each group has one task and each robot can do one task is the online maximum weighted bipartite matching problem (MWBMP). For online MWBMP, it is known that, under some assumptions on the payoffs, a greedy algorithm has a competitive ratio of  $1/3$ . Our key result is to prove that for the general problem, under the same assumptions on the payoff as in MWBMP and an assumption on the number of tasks arising in each group, a repeated auction algorithm, where each group of tasks is (near) optimally allocated to the available group of robots has a *guaranteed competitive ratio*. We also prove that (a) without the assumptions on the payoffs, it is impossible to design an algorithm with any performance guarantee and (b) without the assumption on the task profile, the algorithms that can guarantee a feasible allocation (if one exists) have arbitrarily bad performance in the worst case. Additionally, we present simulation results depicting the average case performance of the repeated greedy auction algorithm

- Luo, L., Chakraborty, N. and Sycara, K. “Distributed Algorithm Design for Multi-robot Generalized Task Assignment Problem”, Proceedings of International Conference on Intelligent Robots and Systems (IROS), Tokyo, Japan, November 3-8, 2013.
- Luo, L., Chakraborty, N., Sycara, K. Distributed Algorithm Design for Multi-Robot Task Assignments with Deadlines for Tasks, International Conference on Robotics and Automation (ICRA), Karlsruhe, Germany, May 6-10, 2013
- Luo, L., Chakraborty, N., Sycara, K., “Competitive Analysis of Repeated Greedy Auction Algorithm for Online Multi-Robot Task Assignment”, International Conference on Robotics and Automation (ICRA), St. Paul, Minnesota, May 14-18, 2012.

## **5. Scalable Human Control of Multi Robot Systems (Lead University of Pittsburgh in collaboration with CMU)**

### **5.1.Overview**

A basic problem in the development of large networked military systems is the integration of humans with unmanned vehicles (UVs). As the number of UVs increases beyond 2 or 3, coordinating their actions becomes too complex for a human operator to manage. Adding additional operators just makes things worse because now each operator must coordinate his UVs with every other operators’ UVs as well as his own. If we allow the UVs to coordinate autonomously the problem becomes trying to find ways to influence their aggregate behavior so they can achieve a range of expected commanders’ intents.


Our work at the University of Pittsburgh and Carnegie Mellon University focused on the problem of human command over multiple UVs. Our objective was to develop techniques that allow human control to scale to increasingly large numbers of UVs. This problem involves both command and monitoring of UVs and effectively exploiting their products.

We have conducted multiple experiments aimed at different aspects of this problem. Here we summarize some of the results.

## **5.2. Control of Multiple UVs Performing Independent Tasks**

### **5.2.1. Teams of Humans Controlling Teams of Robots**

When human operator teams control multiple robots, the way the robots are organized and the methods by which robots are assigned to operators may affect system performance. To check this hypothesis, we completed a large 120 subject study on control of robot teams by teams of human operators. The study addressed the interaction between automation and organization of human teams in controlling large robot teams performing an Urban Search and Rescue (USAR) task. The study used the high fidelity USARSim testbed. Two possible ways to organize operators were identified as individual assignments of robots to operators, *assigned robots*, or as a *shared pool* in which operators serviced robots from the population as needed. The experiment compared two-member teams of operators controlling teams of 12 robots each in the *assigned robots* conditions or sharing control of 24 robots in the *shared pool* conditions.



**Figure 7: USARSim multi robot control system (MrCS) configured for shared control of 24 robots**

An additional comparison was made between *manual* conditions where waypoint control was used, and *autonomy* conditions where autonomous path planning was used

The experiment with teams of two operators replicates the effects of automated path planning found in an earlier single operator experiment (Lewis et al., 2010). In both experiments, relieving operators of the need to perform path planning led to finding more victims and marking their locations more accurately. In the current study participants in the assigned robot condition using automated path planning found twenty-two percent more victims. This gain is particularly significant because the group explored 67% of the map and came close to matching the actual density of victims of .029/m<sup>2</sup>. While the advantages for the autonomy condition are the sort often attributed to situation awareness (SA), process measures suggest the reverse may be true. The times between the appearance of a victim in a robot's camera and marking of the victim were much shorter for the autonomous conditions (*assigned* and *shared pool*). However, the time between selecting a robot to control and marking the victim the robot had found was much shorter in the *manual* conditions as compared to the *autonomy* conditions. In particular, in the *manual* condition we observed times between selecting a robot and marking its victim as low as 14 sec. in the *shared pool* group, approximately one third of the 41 seconds required for *autonomous* operators controlling *assigned* robots. These data suggest that while operators in the *autonomous* path planning condition had more leisure to monitor the cameras leading to earlier detection, once a victim was detected, they had poorer SA for locating the robot and victim on the map.

While team organization was a focus of this study these results were equivocal. *Shared pool* participants across conditions found fewer victims with those controlling *manually* exploring less territory as well. We had hypothesized that increasing automation would improve *shared pool* performance to a greater extent than it improved *assigned* robot performance. This was not seen on any of the measures although the sharp drop off in region explored for *manual* control participants in the *shared pool* condition provided weak evidence for a *shared pool* advantage with *automation*. In the *assigned robot* condition operators on average neglected 2 of their 12 robots, the same number found in (Lewis, et al., 2010). In the *shared pool* condition where robots were not assigned, fewer (8) robots were controlled on average. We attribute this decrement and related effects on team performance to diffusion of responsibility resulting in robots left unattended.

An *unexpected finding* of this experiment was that data from the autonomous conditions did not fit the Neglect Tolerance model well. While the Neglect Tolerance model presumes that human interaction is needed to restore a robot's effectiveness, most interactions in the autonomous version of our task were driven by the detection of a victim rather than degradation of robot performance. We examined the contribution of operators to the system's performance by comparing purely autonomous trials with mixed-initiative ones with operators on hand to provide assistance and found no difference in the regions explored. This leads to new research to refine the Neglect Tolerance model, more precisely define notions of performance in various tasks and construct revised and more realistic theoretical and empirical models.

- Lewis, M., Wang, H., Chien, S., Velagapudi, P., Scerri, P. & Sycara, K. (2010). Choosing autonomy modes for multirobot search, *Human Factors*, 52(2), 225-233.
- Lewis, M., Wang, H., Chien, S., Ma, Z., Velagapudi, P., Scerri, P., & Sycara, K. (2011) Process and performance in human-robot teams, *Journal of Cognitive Engineering and Decision Making*, 5(2), 186-208.
- Lewis, M. & Sycara (2011), Effects of automation on situation awareness in controlling robot teams, *The Fourth International Conference on Advances in Computer-Human Interactions*, 242-248.
- Lee, P., Wang, H., Chien, S., Lewis, M., Scerri, P., Velagapudi, P., Sycara, K. & Kane, B. (2010). Teams for Teams: Performance in Multi-Human/Multi-Robot Teams. *Proceedings of the 54th Annual Meeting Human Factors and Ergonomics Society (HFES'10)*, 438-442.
- Wang, H., Lewis, M., Chien, S., Scerri, P., Velagapudi, P., Sycara, K. & Kane, B. (2010). Teams organization and performance in multi-human/multi-robot teams, 2010 *IEEE International Conference on Systems, Man, and Cybernetics, (SMC'10)*, 1617-1623.

### 5.2.2. Scheduling Human Attention

One case in multi UV control is when UVs perform independent tasks. Under these conditions operators can control UVs in sequence in a round robin fashion. Control of this type resembles a queuing system in which the operator is the server and UV needs for interaction correspond to the jobs. If the operator attention which is shifted from UV to UV can be more effectively scheduled using Operations Research techniques then system performance could be improved. There are two prerequisites for this approach: 1) demonstrating that human attention can be effectively scheduled and 2) developing formal scheduling models offering improvement for multi UV control (see section 4).

We have conducted an experiment investigating the effectiveness of directed attention and open alarming for improving operator response to UV failures in a multi UV system. This work is reported in:

- Chien, S., Mehrotra, S., Brooks, N., Sycara, K., & Lewis, M. (2011) Effects of Alarms on Control of Robot Teams, *Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11)*.

Motivated by these results we conducted a series of experiments to see how our test environments could be made more failure-prone in order to require more human intervention and how we could alert the operator to failure detected through self-reflection. These pilots have led to the redesign of our test environment making the tasks more difficult by reducing lighting, adding smoke and debris. We have also equipped our simulator with the capability of injecting failures so arrival rates for tasks demanding operator attention can be controlled allowing us to more closely match queuing models and test approaches to operator aiding. An experimental study found advantages for alerting operators to failures but not for directing them in a sequence for addressing the failures. In a follow-on experiment we found that where there were substantial

advantages to a particular sequence of interactions (shortest job first, SJF, discipline) performance could be improved by directing operators to this sequence.

Chien, S., Mehrotra, S., Brooks, N., Lewis, M. & Sycara, K. (2012). Scheduling Operator Attention for Multi-Robot Control. Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2012).

### **5.2.3. Human Operator Utilization**

Operator utilization refers to the proportion of time an operator is performing a task. Studies have shown that over a wide range of settings performance deteriorates at utilizations above 75%. We have developed a synthetic air traffic control task that allows us to control operator utilization precisely and to score each action for latency and correctness while controlling for difficulty. We run an experiment comparing aggregations of work-rest intervals of varying lengths. Results were reported in:

Lee, P., Kolling, A., & Lewis, M. Workload Modeling using Time Windows and Utilization in an Air Traffic Control Task, *Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society* (HFES'11).

Lee, P., Kolling, A. and Lewis, M. Combining latency and utilization in investigating human operator workload, *2011 IEEE International Conference on Systems, Man, and Cybernetics*, (SMC'11)

### **5.3.Human Search Using Algorithmically Generated Paths**

Humans use a variety of information about the environment in planning paths and typically generate relatively straight paths with few turns or backtracking. Automated path planners by contrast, rely strongly on local data and as a consequence generate less smooth paths. We conducted experiments to see whether operators could perform as well at a search and rescue task using algorithmically generated paths as with those generated by another human. These results were reported in:

Chien, S., Wang, H., & Lewis, M. (2010). Human vs. algorithmic path planning for search and rescue by robot teams, *Proceedings of the 54th Annual Meeting of the Human Factors and Ergonomics Society* (HFES'10), 379-383.

Scerri, P., Velagapudi, P., Sycara, K., Wang, H., Chien, S. & Lewis, M. (2010). Towards an understanding of the impact of autonomous path planning on victim search in USAR, *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS'10), 383-388.

## **6. Development of queuing models to characterize and aid multi robot control (Lead: CMU-Robotics)**

### **6.1.Service level Differentiation**

We explored the effects of service level differentiation on a multi-robot control system. We investigated the conjecture that duration of human interaction, interaction time (IT),



and quality are correlated with performance and length of the subsequent neglect interval, called neglect time (NT) and explored the tradeoffs for multirobot systems. We examined the premise that although long interaction time between robots and operators hurts the efficiency of the system, it allows robots longer neglect times and better performance thus benefiting the system. We addressed the problem of how to choose the optimal service level for an operator in a system through a service level differentiation model. The model identifies the optimal service strategy to maximize system performance in multi-robot control through a service level differentiation method based on two types of service: high-quality-long-time and low-quality-short-time. The operator offers high quality service with probability  $p$  and low quality service with probability  $1-p$ . The problem is to find the probability  $p^*$  that maximizes system performance.

Modeling different levels of service is motivated by real human performance data which shows a wide variety of ITs related to variations in demands on the operator. While the earlier neglect tolerance model assumed a fixed efficiency threshold for each robot our model relates IT and NT to optimal system performance allowing the individual thresholds to vary. This increased flexibility not only improves team performance but agrees with human data showing performance per robot to decrease smoothly with increasing team size rather than dropping abruptly upon reaching the fan-out threshold. We modeled service level differentiation in two types of queuing systems (a) open queue, and (b) closed queue. Open queue systems make the assumption that robots arrive at the queue according to some arrival process (usually Poisson), get serviced and then leave the system. Most of queuing models in the literature are open queue because they are easier to analyze. We were able to find exact analytic solution for the optimal  $p^*$  in the open queue model of service differentiation. This is an important contribution.

While an open system model may provide an approximation of systems with long NTs it is limited in its ability to accommodate the assumption of repeated interactions made by the neglect tolerance model and fan-out estimators. To address this, we developed the *first closed system model* for human-robot teams that meets the assumptions of Crandall's (2005) informal neglect tolerance model. A closed system model is one where robots arrive, get served and *return for service*. Close queue models are far more difficult to construct and analyze than open queue ones because of the interdependence between the service process and the arrival process. Close queue models are even more challenging to develop and analyze when service differentiation is also modeled. However, close queue models with service differentiation are applicable to human control of multiple robots since typically the operator controls a known number of robots that may require repeated service during system operation, thus returning to the queue.

Since it is extremely challenging to find exact analytic optimal solutions for close queue models we developed techniques to find solutions algorithmically. Experimental results comparing system performance for different values of system parameters show that a mixed strategy is a general way to get optimal system performance for a large variety of

system parameter settings (e.g.; different values of  $\lambda$ , the arrival rate parameter of the Poisson process, number of robots etc) and in all cases is no worse than a pure strategy. Results were reported in:

Xu, Y, Dai, T., Sycara, K. & Lewis, M. (2010). Service level differentiation in multi-robots control, *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10)*, October 18-22, Taipei, Taiwan, 2224-2230.

## 6.2 Game Theoretic Model of Queuing to Schedule Operator Attention

In order to increase human span of control, increased robot automation is needed. In particular, the ability of robots to self-reflect and self-monitor frees the operator from having to monitor the progress of the robots. This, in turn increases the neglect time, given a particular interaction time. We developed a game-theoretic queuing model that addresses robot self-assessment in human-robot-interaction systems. Four issues were incorporated based on the theory of queuing and performance: 1) individual differences in operator skills/capabilities, 2) differences in difficulty of presenting tasks, 3) trade-off between human interaction and performance and 4) the impact of task heterogeneity in the optimal service decision-making and system efficiency. Our model makes the additional plausible assumption that increasing the human operators' skill level or the service duration (interaction time) will lead to equivalent or longer subsequent neglect times. We explore the situation in which UVs are empowered with self-assessment and can choose their operator rather than requiring a centralized queue manager.

Our model takes into account a variety of parameters likely to affect multi UV control. The single-human/multi-robot system is modeled as an open queuing system in which different types of arriving UVs require varying degrees of attention (reservation utility) with differing costs of continuing to operate in their degraded mode (waiting costs). Our key findings include:




Fig. 3. The sensitivity of optimal service rate and system utility in  $c$ . Assume  $V_b = 0.3, \mu_b = 2, \alpha = 3, r = 0.8$ .




Fig. 7. The sensitivity of optimal service rate and system utility in  $r$  considering task heterogeneity in waiting costs. Assume  $V_b = 0.3, \mu_b = 2, \alpha = 3$ .

**Figure 8: Two figures showing sensitivity of optimal service rate to performance under various conditions**

- In the baseline model, all the robot tasks are assumed to be homogeneous in both reservation utility and waiting costs. The optimal service rate is shown to be increasing in the human operator's skill level and decreasing in the reservation utility. Counter-intuitively, we also show that the optimal service rate decreases in the waiting costs (see Figure 8 right hand figure). In other words, the more impatient each robot is, the more time the human operator should spend on servicing it. The rationale is that the human operator provides *value-added service*, and higher service quality is required to compensate for utility loss associated with queuing time.
- When task heterogeneity in waiting costs is incorporated, we show that the optimal service rate still increases in the human operator's skill level. However, an increased reservation utility can lead to either a higher or a lower optimal service rate (see Figure 7).
- When the task heterogeneity in reservation utility is accounted, we show that the optimal service rate increases and stays roughly constant as the waiting costs increases. This is different from the baseline model since in this case a higher waiting cost increases the system's pressure for speeding up and reducing the system delay.

The simplicity of our model allows it to be extended to more complex situations and can be easily used in applications. We have also investigated the multi-operator-multi-robots case in which the tradeoff lies not only in the one-shot interaction between robots and a single operator, but also in how to coordinate different human operators so as to achieve the best system performance.

Work developing scheduling models for improving performance of human multi UV systems was reported in:

Dai, T, Sycara, K., Lewis, M. A game theoretic queuing approach to self-assessment in human-robot interaction systems. IEEE International Conference on Robotics and Automation (ICRA 2011), May 9-13, Shanghai, China, 2011.

Ying Xu, Tinglong Dai, Katia Sycara, Michael Lewis. 2012. A Mechanism Design Model in Multi-Robot Service Queues with Strategic Operators and Asymmetric Information. Proceedings of the 51st IEEE Conference on Decision and Control: CDC'12

## **7. Scalable Displays (Lead: U of Pittsburgh in collaboration with CMU)**

A complementary approach to using autonomous coordination of robots in order to increase the operator's span of control, is to (a) reduce the operator's burden of monitoring the UV cameras and (b) helping in managing the vast amounts of information coming from the cameras. To help with reducing the operator's monitoring burden, we developed techniques to allow robot self-reflection. Self-reflection allows the robots to report suspected failures, thus alleviating operator monitoring for failures, though the primary monitoring task of searching for victims is still left to the operator. The queuing

approaches (see section 6) that we have developed and tested allow the self-reporting robots to appear as customers in a queue, thus allowing optimized scheduling of operator attention. To allow the operator to best manage the amounts of information returned from the cameras, we have developed asynchronous display approaches that allow the operator to inspect *non redundant imagery* in context.

### **7.1 SUAVE**


The problem is simplest for UAV images which can be textured onto a map. New images of a location replace old ones and the map provides a spatial context for the images. Earlier picture-in-picture displays used the approach of painting imagery onto a map to provide context, however, as an asynchronous display, SUAVE allows the operator to inspect the entire map using world-in-miniature and fly-through techniques. Experiments testing this approach were reported in:

Abedin, S., Brooks, N., Owens, S., Scerri, P., Lewis, M., & Sycara, K. SUAVE: Integrating UAV Video Using a 3D Model, *Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11)*.

Abedin, S., Wang, H., Lee, P., Lewis, M., Brooks, N., Owens, S., Scerri, P. and Sycara, K. SUAVE: Integrating UAV Video Using a 3D Model *2011 IEEE International Conference on Systems, Man, and Cybernetics, (SMC'11)*

### **7.2 Image Queue**

Organizing UGV imagery is more difficult than for UAVs because it has no natural organizational context such as a map. The same object will show great variation in size and appearance as it is viewed from different angles and distances. When multiple UGVs are involved it can be extremely difficult sorting out camera views to identify overlaps. Our experimental Image Queue display addresses this problem by storing video along with UGV pose and location. The database is then searched to identify images providing the greatest additional visual coverage. This has required a more sophisticated search in which visual coverage is coordinated with mapping. During the search the operator examines a small number of prioritized “film strips” to see what has been seen by the team of robots. By assembling a collection of non-overlapping high coverage images, the



**Figure 9: Image Queue selects non-redundant images: image 1 is selected since image 2 is contained wholly within image 1 and image 3 is contained partially within image 1**

display allows the operator to observe most of the information contained in large pool of imagery collected by a UGV team. The first experiment compared search and rescue performance between operators using the image queue and others relying on streaming video. In a second experiment the utility associated with gains in coverage was augmented with ATR for victims in selecting imagery to be viewed. In the current test environment after a search is complete, the top ten frames in the queue account for more than 70% of the map while the top 100 account for over 99%.

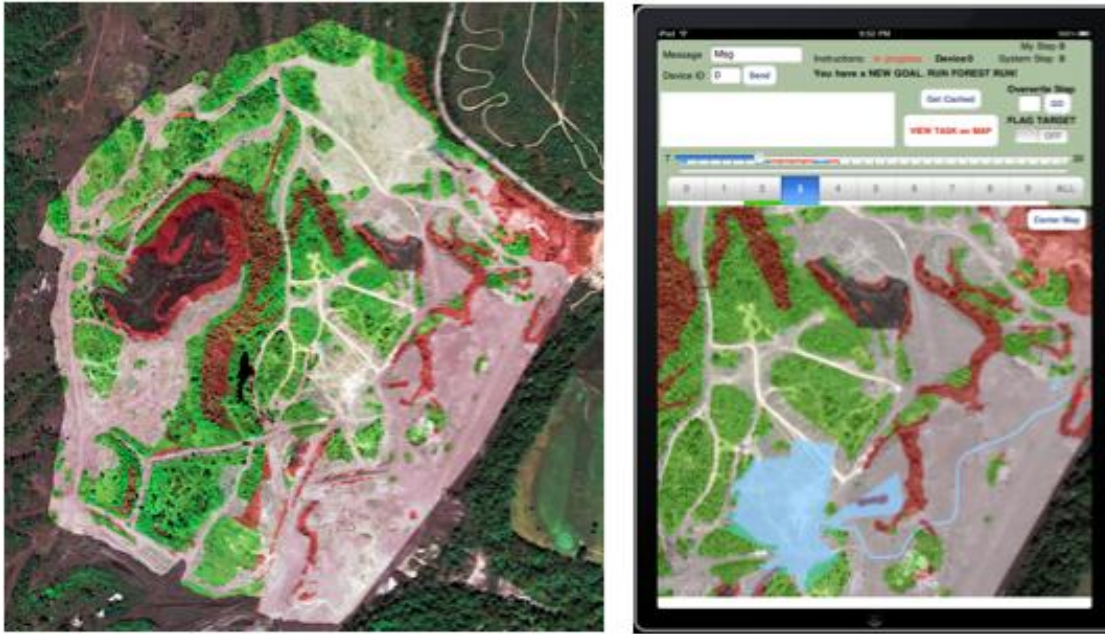
Results were reported in:

- Wang, H., Kolling, A., Abedin, S., Lee, P., Chien, S., Lewis, M., Brooks, N., Owens, S., Scerri, P. & Sycara, K. (2011) Scalable target detection for large robot teams, *Proceedings of the 6th ACM/IEEE International Conference on Human-Robot Interaction*.
- Scerri, P., Owens, S., Sycara, K. & Lewis, M. (2010). User evaluation of a GUI for controlling an autonomous persistent surveillance team, In *SPIE'10*.
- Brooks, N., Wang, H., Chien, S., Lewis, M., Scerri, P., & Sycara, K. Asynchronous Control with ATR for Large Robot Teams, *Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11)*.
- Wang, H., Chien, S., Lewis, M., Brooks, N. and Sycara, K. Image Queue: Scalable Display for Multiple Robots, *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011)*.

## **8. Dynamic Targets (Lead: U of Pittsburgh in collaboration with CMU)**

We developed an approach for a pursuit-evasion problem that considers a 2.5d environment represented by a height map. Such a representation is particularly suitable for large-scale outdoor pursuit-evasion, captures some aspects of 3d visibility and can include target heights. In our approach we constructed a graph representation of the environment by sampling strategic locations and computing their detection sets, an extended notion of visibility. From the graph we computed strategies using previous work on graph-searching. These strategies were used to coordinate the robot team and to generate paths for all robots using an appropriate classification of the terrain. In experiments we investigated the performance of our approach and provided examples including a sample map with multiple loops and elevation plateaus and two realistic maps, a village and a mountain range. To the best of our knowledge the presented approach was the first viable solution to 2.5d pursuit-evasion with height maps.

To examine whether the approach would be useful in realistic environments, we conducted a pilot experiment with 10 humans, 8 pursuers and 2 evaders. The environment was Gascola, a wooded and uneven terrain area outside of Pittsburgh. The 2 evaders were free to move as they pleased to avoid detection; the 8 pursuers, each carrying an iPad, acted as robots, obeying the directions of the algorithm, given to them via a GUI. (see figure below). The pursuers were successful in all trials.



**Figure 10: (a) Satellite map of Gascola overlaid with a mask denoting nontraversable terrain (red), shrubs and trees (green). Darker areas are not part of the experiment while lighter areas are. A graph is overlaid on the map (not shown here) that allows generation of best locations and paths for the pursuers to follow. (b) Screenshot of the iPad application showing satellite imagery and the mask. Agents are instructed to go to goal locations and receive a suggested path shown with a light blue line. The area an agent is responsible for at a step is marked with a light blue polygon. (a) (b)**

Results of these experiments were reported in:

Kleiner, A., Kolling, A., Lewis, M., Sycara, K., “Hierarchical visibility for guaranteed search in large-scale outdoor terrain“, *Journal of Autonomous Agents and Multi-Agent Systems*, 2011, DOI 10.1007/s10458-011-9180-7

Kolling, A., Kleiner, A., Lewis, M. Sycara, K. Computing and executing strategies for multi-robot search. IEEE International Conference on Robotics and Automation (ICRA 2011), May 9-13, Shanghai, China, 2011.

Kolling, A., Kleiner, A., Lewis, M., & Sycara, M. (2010). Pursuit-evasion in 2.5d based on team-visibility, Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS’10), 4610 – 4616.

## **9. Human Influence of Robotic Swarms (Lead: U. of Pitt in collaboration with CMU)**

### **9.1. Human Control of Swarms**

Many approaches to coordinating large numbers of UVs rely on local control laws and emergent behavior. Because behavior is emergent rather than designed a priori it is difficult to define mechanisms allowing human control. We have begun systematic

research on this problem using a limited number of “communication graphs” that constrain behavior to maintain connectivity and seeking ways through manipulation of connectivity and basic coordination algorithms (rendezvous, deployment, boundary following) to allow human control

Behaviors of swarm robotic systems can be influenced by a human by altering the behavior of some swarm members, altering the control laws that the individual swarm members use or altering the environment in which the swarm operates. We have systematically investigated the effect of influencing the swarm through these three different schemes. Our research efforts were geared towards understanding the following key questions: (1) For swarm robotic systems when does human influence benefit the overall system? (2) What type of influence, namely, directly influencing swarm member behaviors or influencing swarm behaviors through environment modification helps human operators perform better, if at all? (3) How does the mismatch in operator understanding of swarm state and swarm member understanding of operator intent affect the performance of the overall system? (4) How can the adverse effects of operator-swarm state or intent mismatch be mitigated?

To answer the above questions, we have conducted theoretical studies as well as human-subject experiments, which we believe are the first of its kind in the context of human control of swarm robotic systems. For the experiments we used the task domain of information foraging. Our key findings are as follows:

- We find that although the autonomous algorithms perform better than humans in very simple environments (with no obstacles) as the complexity of the environment increases, the human performs significantly better. We also find that novice human operators perform better by directly influencing the swarm members rather than by altering the environment (by activating/deactivating beacons in the environment).
- When there is an intent mismatch between the human and the robots and the operator is unaware of the exact state positions (either due to limitations in communication bandwidth or communication delay), the operator performance decreases significantly when compared to the complete state information condition. However, using techniques like display of statistics of the spatial distribution of the agents or predictive display, it is possible to mitigate the effects of uncertainty.
- We have introduced the concept of *neglect benevolence*, as a “meta-strategy” that human operators use to control swarms. In neglect benevolence, a human operator allows the swarm to evolve on its own before giving new commands. From our experiments, we find that operators exploited neglect benevolence in different ways to develop successful strategies for controlling the swarm in the presence of uncertain information about the swarm state.

## **9.2. Principles of human control for large swarms**

Our first experiment was set up for investigating principles of human control for large swarms. The swarm robots were given some coordinated behaviors like deployment, flocking, and rendezvous along with some primitive behaviors like stop, go to a target

location, and random movement. The operator could control the robots in two ways (a) by selecting a subset of robots and specifying a behavior for them (called *selection control* hereafter) and (b) by placing beacons in the environment (called *beacon control* hereafter) that could set the behavior mode of robots within a certain distance of the beacon (this is a way of influencing robots by “modifying the environment”). We chose five different environments of different levels of complexity (see Figure ).

The results indicate that in environment (a) and (b) the autonomous algorithm performs significantly better than the human operator using selection control. However, in more complex environments (maps (c), (d), and (e)), selection control by the human leads to significantly better performance. On the other hand, using beacon control, the operators either underperformed compared to the other two conditions or there was no significant difference in performance. Furthermore, as the number of robots increased, although both selection control and beacon control showed decrease in performance, the decrease for selection control was much less than that for beacon control. Thus, these findings seem to suggest that it is easier for novice operators to control the swarm robots directly rather than through the environment.

Walker, P., Kolling, A. and Lewis, M. Human Exploration Patterns in Unknown, Time-sensitive Environments, *2011 IEEE International Conference on Systems, Man, and Cybernetics*, (SMC'11)

Walker, P., Amnipur Amraii, S., Lewis, M., Chakraborty, N., Sycara, K. Human Control of Leader-Based Swarms *Proceedings of the Conference on System Man and Cybernetics, Manchester, UK., October 13-16 2013.*

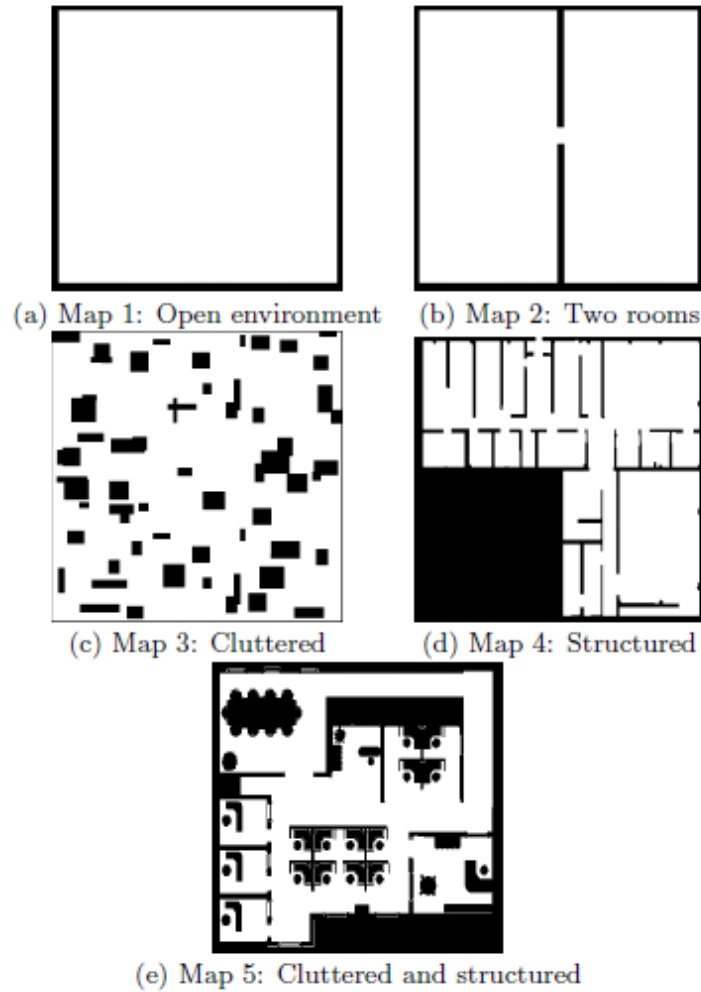
### **9.3 Swarm Control with imperfect information**

One assumption that is implicitly made in most of the work on human control of swarms is that the swarm “understands” the operator intent perfectly and the operator knows the state (e.g., positions of the swarm members perfectly), i.e., the operator is omniscient. However, in practice, these assumptions are often violated. Two key challenges in human swarm interaction are that (a) the state information of the robot available to the human may not be accurate and (b) there may be a mismatch between the intent of the operator and the robots understanding of the human intent. The error in the swarm state available to the human and the intent mismatch can happen due to communication limitations (e.g., bandwidth limitations or communication latency) and localization error of individual robots. We performed experimental studies focusing on the effect of communication bandwidth limitations and communication latency on human control of swarms.

#### **9.3.1 Swarm Control with Communication Bandwidth Constraints**

Limited communication bandwidth is a constraint that arises in many practical scenarios such as undersea missions or networks of limited capability robots. In our experimental scenario, a human operator has to guide a robotic swarm to find unknown targets in a given area. The area is divided into a finite number of regions (whose boundaries are unknown to the interface) and the operator has to match the target found to the regions. The robots have a single behavior, namely achieving consensus on direction on motion. The humans can guide the swarm by giving them a point in the environment towards






**Figure 11 : Five tests environment of different complexity. Obstacles are black and free spaces are white.**

which the robots have to travel. The robots are assumed to have a localization error and the robot position and orientation is assumed to be a Gaussian distribution.

In our experiment each subject performs the mission under three conditions (that are presented to them in a random order), namely, (a) low swarm-to-human bandwidth and low intra-swarm bandwidth (low bandwidth condition), (b) low swarm-to-human bandwidth and high intra-swarm bandwidth (medium bandwidth condition) and (c) high bandwidth between swarm and operator (high bandwidth condition). For low bandwidth condition, we assume that only one robot can send its state information at a time instant, this assumption creates displayed information that lacks temporal and spatial resolution. For the medium bandwidth condition, the swarm communicates among themselves to estimate their mean orientation and standard deviation of orientation, which is displayed on the screen creating a limited spatial resolution of the swarm's state. In the high bandwidth condition, all the robots could send their position and orientation information

to the operator creating high spatial and temporal resolution given the errors of the individual robots.

Our experimental results (see Figure (12)) indicate that, as expected, there is a degradation of performance in the low bandwidth condition compared to the high bandwidth condition. However, in the medium bandwidth condition, where the *human had an understanding of the state of consensus of the robots (and thereby whether the robots were moving in the direction the human desired) from the standard deviation of orientation*, they performed as well as the high bandwidth condition. These results show that even in the absence of complete information about the swarm states, if task-appropriate statistics of the swarm is displayed to the user, the effects of incomplete state information can be mitigated. Results are reported in selected publications below.



**Figure 12 : Performance of the medium and high bandwidth condition is comparable, while both outperform the low bandwidth condition**

Kolling, A., Sycara, K., Nunnally, S., Lewis, M. Human Swarm Interaction: An Experimental Study of Two Types of Interaction with Foraging Swarms, *Journal of Human-Robot Interaction*, June 2013.


Nunnally, S., Walker, P., Lewis, M., Kolling, A., Chakraborty, N., Sycara, K. & Goodrich, M. Human Influence of Robotic Swarms with Bandwidth and Localization Issues, *2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC'12)*, Oct 14-17, Seoul, Korea, 2012

### 9.3.2 Swarm Control with Communication Latency

A second experiment investigated effects of communication delay on human performance in controlling swarms. In many operational settings, human operators are remotely located and the communication environment is harsh. Hence, there exists some latency in information (or control command) transfer between the human and the swarm. In our experimental foraging scenario, a human operator guides a swarm to find unknown targets in a given area. The robots have a single behavior, namely flocking, and the operator applies inputs (a) to give a desired direction of flocking to the robots and (b) to enforce cohesiveness among the robots (by activating constraints for attracting neighbors that are far away and repelling neighbors that are very close). In our experiment, each subject performs the mission under three conditions, namely, (a) without any latency

(control condition), (b) with equal latency in the human to swarm and swarm to human communication channel (c) the same latency as (b) but with a predictive display. In all conditions, each robot has some error in transforming the orientation heading to its own reference frame (due to localization errors), which is modeled as a Gaussian distribution.

Our experimental results (see Figure (13)) indicate that, as expected, there is a degradation of performance due to latency. However, when *using the predictive display, the performance of the operators can be as good as it was in the absence of delay* (control condition). We also found that the users exhibited different strategies for effectively controlling the swarm.



**Figure 13: The performance of the operators with latency and predictive display is comparable with the control condition of no latency and significantly better than the latency condition without predictive display.**

### 9.3.3 Neglect Benevolence

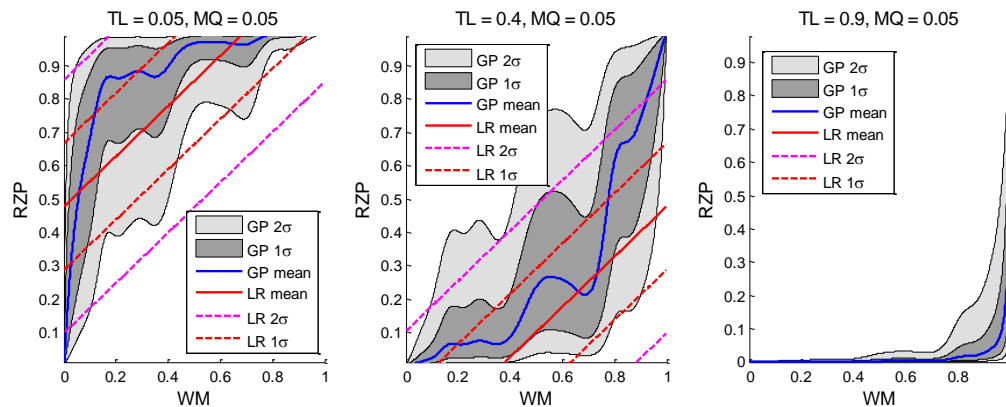
The human operator needs to influence the swarm without adversely disturbing the swarm (such as breaking it into many small connected components). The effect of an operator command is dependent on swarm state, which gradually evolves to a steady state after a command has been issued. To capture the idea that humans may need to observe the evolution of the swarm state before acting, we investigate a novel concept called *neglect benevolence*, whereby neglecting the swarm before issuing new commands may be beneficial to overall mission performance. Our results show that operators came up with different strategies by exploiting neglect benevolence that resulted in improved performance. In general, human operators are limited in their ability to estimate the best time to give input to the swarm, (e.g. when mission goals change). Therefore, automated aids that calculate the optimal input time could help the human operator achieve best system performance. This raises the important question of the existence and means of calculation of the optimal time for the operator to give input to the swarm in order to optimize swarm behavior. This could have significant practical implications. Therefore, we (a) formally defined the new notion of Neglect Benevolence, (b) we proved the existence of Neglect Benevolence for a set of linear dynamical systems, (c) , we provided an analytic characterization and an algorithm for calculating the optimal input time.

Walker, P., Kolling, A., Chakraborty, N., Nunnally, S., Sycara, K. & Lewis, M.. Neglect Benevolence in Human Control of Swarms in the Presence of Latency, 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC'12), Oct 14-17, Seoul, Korea, 2012.

## 10. Human Decision Making in the Presence of Complex Automation (Lead: Cornell)

### 10.1. Modeling of Human Decisions in complex DDD games. (in collaboration with GMU)

It is becoming increasingly important to be able to predict human operator effectiveness and performance in large scale human-automation systems. A key problem is to determine how human performance in such systems changes under varying task conditions in applications that prohibit exhaustive experimental evaluation. To this end, relevant tasking conditions and cognitive factors such as working memory can be used to construct scalable probabilistic human performance models from limited experimental data. Our groups studied different statistical modeling methods for predicting human operator performance in a DDD air defense simulation scenario, where several performance metrics were modeled as a function of task load, message quality, and operator working memory capacity. It was found that state-of-the-art Gaussian Process (GP) regression models can make predictions with uncertainty bounds that are as good as or better than simple linear regression and discrete Bayesian network (BN) prediction models. While the probabilistic nature of GP and BN models was found to be very useful in removing irrelevant/unimportant factors for predicting certain performance measures, these models also demanded more computational resources for learning and



**Figure 14: Predicted Red Zone Safety Performance (RZP) for DDD experiments. RZP mean and standard deviations for GP and simple linear regression (LR) models using novel input values for task load (TL), message quality (MQ), and working memory (WM) values not observed in training data. Note that LR results are completely negative in the last plot. TL values from left to right correspond to scenarios with 10, 80, and 180 enemy aircraft, respectively (note that only 31 or 47 enemy aircraft were actually encountered in the experimental trials).**

implementation than simple linear regression. As such, the usefulness of each model in predicting human performance depends strongly on computational constraints and the availability of experimental data for a particular application.

N. Ahmed, M. Campbell, "On Estimating Simple Probabilistic Discriminative Models with Subclasses," *Expert Systems With Applications*, published on-line Dec 2011, Vol 39, No 7, June 2012, pp 6659–6664.

## **10.2. Human-Robotic Information Fusion**

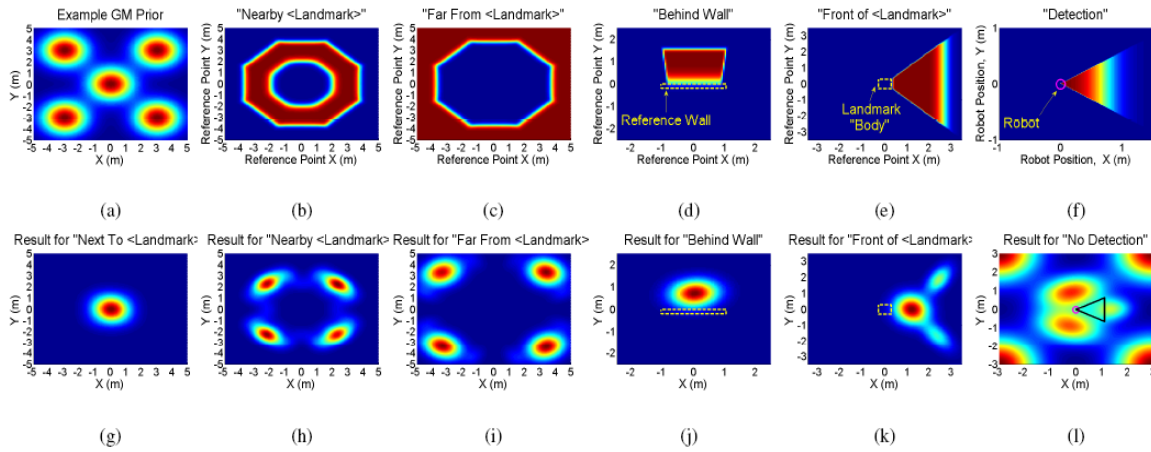
Although humans play important roles as both operators/supervisors in human-robot systems, their ability to contribute useful information (beyond object classification) in many scenarios has been largely overlooked. Given the limited amount of information obtainable through robot perception alone, proper fusion human-generated information could greatly enhance the situational awareness and performance of human-robot teams in applications such as surveillance and target search. Since humans tend to compress information about various physical phenomena into "fuzzy" discrete categories when relating observations, appropriate human "likelihood models" can be modeled probabilistically via machine learning techniques. "Soft" observations under these likelihoods can then be recursively fused with conventional "hard" robot sensor data in a rigorous Bayesian manner, so that human agents can be treated as soft sensory input channels. We developed a sensor fusion architecture, where robots and humans can fuse information at different levels of the perceived model, as would intuitively occur because humans are, in general, good reasoners.

We have developed recursive data fusion approximations for a wide class of soft human sensor observations using variational Bayes, importance sampling, and Gaussian mixture modeling techniques. Furthermore, we have experimentally validated the proposed fusion strategy on a real multi-target search problem with a human-robot team. The approach uses a Bayesian estimation framework for mapping and classifying objects in the surrounding of a mobile robot based on 2D laser range data and additional human input. Object observations made by humans through the robot's camera are treated as additional probabilistic observations inside a recursive Bayes estimator for determining an object's ID. A Rao-Blackwellized particle filter implementation is chosen for simultaneously estimating the locations of objects, location measurement to object associations, and object class associations. Reliably detecting and identifying objects is one of the necessary basic skills of service robots, and this problem is far from being solved.

Our results show that the proposed recursive Bayesian fusion of human and robot information leads to superior search performance in terms of mission completion time and number of targets found, even with poor prior target information.

Ahmed, N. and Campbell, M., "Variational Bayesian Learning of Probabilistic Discriminative Models with Latent Softmax Variables," *IEEE Transactions on Signal Processing*, on-line April, 2011, Vol 59, No 7, July 2011, pp 3143-3154

- R. Tse, N. Ahmed, M. Campbell, "Unified Mixture-Model Based Terrain Estimation with Markov Random Fields," 2012 IEEE International Conference on Multisensor Fusion and Integration.
- N. Ahmed, J. Schoenberg, M. Campbell, "Fast Weighted Exponential Product Rules for Robust General Multi-Robot Data Fusion," *Robotics Science and Systems Conference*, 2012.
- E. Sample, N. Ahmed, M. Campbell, "An Experimental Evaluation of Bayesian Soft Human Sensor Fusion in Robotic Systems," 2012 *AIAA Guidance, Navigation and Control Conference*.



**Figure 15 - (a) Gaussian mixture prior pdf for target location. (b)-(e) Likelihood models for soft human observations. (f) Likelihood for robot's visual target detector. (g)-(l) Posterior Gaussian mixture pdfs resulting from Bayesian fusion of corresponding observations in (b)-(f).**

### **10.3. Integration of Perception and Planning in Human-Vehicle systems, (in collaboration with MIT)**

Large scale human-robot teams must carefully coordinate their efforts to complete multiple tasks in highly uncertain environments. In particular, the need to gather more information to reduce uncertainty in such environments must be balanced with the need to complete all required tasks in an efficient and timely manner. The goal of this work is to develop robust probabilistic methods for sharing tasks and all available information relevant to those tasks among a networked team of multiple human-robot agents. This work focuses on three key aspects of networked human-robot team cooperation in uncertain environments: decentralized high-level information-based task planning; local information-based low-level task execution; and Bayesian fusion of robot sensor data with observations obtained from human agents. Hardware based experiments based on an indoor multi-target search application with an actual human-robot team were conducted to assess the performance of Consensus Based Bundle Adjustment (CBBA) algorithms task allocation with two different task execution strategies (IRRT vs. Greedy MDP path planning) and human/robot data fusion modalities (robot data only vs. robot + human data). The results show that it is possible to greatly enhance human-robot team performance (e.g. in terms of number of targets found, time to find all targets, and distance traveled by robots) with the proposed planning strategies as long as they are tuned appropriately to handle spontaneous human information reports.

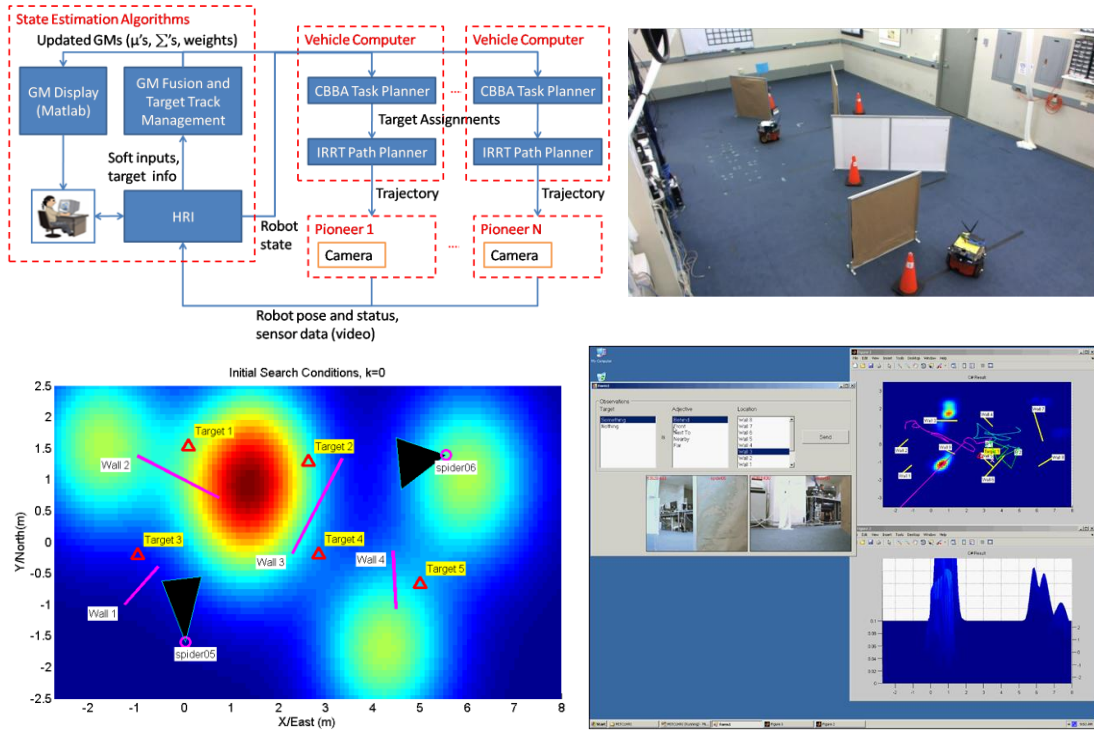


Figure 16: MIT-Cornell collaboration: Real-time information-rich task allocation, trajectory planning and target estimation for human-robot search and track missions. Images show: real-time fusion architecture (top-left), Gaussian multi-modal fusion for target estimation (top-right), real-time experimental search and track mission using human-robot team (bottom-left), and human-robot interface (HRI) depicting operator's view and interface options for soft inputs (bottom-right)




Figure 17: – Probability of locating each of 5 targets at their true locations over time for indoor multi-target search experiment with human-robot team. Probabilities are calculated using Gaussian mixture pdfs that represent uncertainty in target locations over the search map following data fusion. Left plot shows probabilities when only robot sensor data is fused together; the robots are not confident that the targets are near their true locations and so take longer to find them using a greedy search. Right plot shows probabilities when human observations are fused with robot data; the robots become more confident that targets are near their true locations and take less time to find the targets. In both plots, the robots are re-assigned targets to search every 4 sec by CBBA.

Results were reported in:

Ponda, S., Ahmed, N., Luders, B., Sample, E., Levine, D., Hoossainy, T., Shah, D., Campbell, M., and How, J. P., “Decentralized Information-Rich Planning and Hybrid Sensor Fusion for Uncertainty Reduction in Human-Robot Missions,” *AIAA Guidance, Navigation and Control Conference*, (GNC), Portland, OR, August 2011 (**Best paper award**).

#### **10.4 Human Network Experiments**

We performed an experiment where five humans were networked together and used handheld PCs to perform a search experiment outdoors. Ad hoc networking was performed using handheld computers; uncertain variables fused included 1) yes/no found target; 2) human location and motion (GPS with a motion filter); 3) human head orientation (for looking); 4) uncertainty model in human’s ability to find target (bearing and range), found empirically from human decision data. A key element was exploring different sharing methodologies which maintained probabilistic formalism, yet could be implemented on computers. Interesting elements of the experiments came out, such as fusing very uncertain data as people walked through buildings. We completed a series of experiments with the five human nodes.

#### **10.5. Qualitative Path Planner**

We developed formal inference algorithms that enable humans to qualitatively draw plans that robots can then follow. The Qualitative Path Planner (QPP) is a proposed method for controlling a mobile robot using qualitative inputs in the context of an approximate map, such as one sketched by a human. By defining the desired trajectory with respect to observable landmarks, human operators can send semi-autonomous robots into areas for which a true map is not available and teleoperation is not desirable. Such applications may include planetary exploration, in which large communication delays necessitate more autonomous navigation while still keeping the human operator ‘in charge’ of the robot, or military/rescue operations that may require teams of robots to operate in unmapped environments or areas with poor communication.

### **11. Human Team Interaction with Automation (Lead: GMU)**

#### **11.1. Linear and Bayesian Probabilistic Models of Networked Human System Performance**

Human-automation performance in a dynamic decision-making task requiring supervision of multiple unmanned air vehicle (UAV) assets was examined and modeled, in two parts. First, a human-in-the-loop simulation experiment was carried out examining human-UAV system performance under different levels of task load that posed increasing demands on the operator’s working memory capacity (de Visser et al., 2010). The effects of a networked environment on performance were also examined by manipulating the number and quality of network message traffic to the human operator provided by an



automated agent. Both task load and message quality affected performance, but these effects were modulated by individual differences in participant working memory capacity. The performance data were then analyzed using linear regression and Bayesian probabilistic models namely Bayesian networks and Gaussian processes (Ahmed et al., 2011). Working memory capacity was a parameter in all the models. The relative utilities of the different models in prediction of several different aspects of human-automation performance were evaluated. While linear regression and Gaussian processes provided the best overall predictions, the “best” model for a specific application depends on desired tradeoffs between computational complexity, performance requirements, and data availability.

Data were obtained for the effects of different levels of task load and network message quality on human-UAV system performance. Both task load and message quality affected human-automation performance, but these effects were modulated by individual differences in participant working memory capacity. These data were used to learn predictive statistical operator performance models based on classical linear regression, probabilistic Bayesian networks (BN), and nonparametric Gaussian processes (GPs), where individual operator working memory capacity was a parameter in all models. The linear and GP performance models provided the best overall predictions, while the BN and GP models were most robust to the influence of irrelevant factors. The results support the conclusion that high inter-individual variability can be dealt with by including operator working memory capacity in all such statistical models. However, the “best” model for a specific application depends on desired tradeoffs between computational complexity, performance requirements, and data availability. Finally, the GP models also allowed for prediction of performance in cases where experimental data were not available (e.g., larger number of UAVs, greater network message complexity, higher operator working memory capacity). If validated in follow-up analyses, these models will achieve one of the overall goals of the MURI project, namely “scaling up” of models of networked human-automation performance.

Ahmed, N., de Visser, E., Shaw, T., Mohamed-Ameen, A., Campbell, M. A., & Parasuraman, R. (2012). Predicting human-automation performance in networked systems using statistical models: The role of working memory capacity. *Interacting with Computers* (in press).

Ahmed, N. and Campbell, M., “Variational Bayesian Learning of Probabilistic Discriminative Models with Latent Softmax Variables,” *IEEE Transactions on Signal Processing*, Vol 59, No 7, July 2011, pp 3143-3154.

de Visser, E., Shaw, T., Mohamed-Ameen, A., & Parasuraman, R. (2010,). Modeling human-automation team performance in networked systems: Individual differences in working memory count. Proceedings of the Human Factors and Ergonomics Society, Santa Monica, CA: Human Factors and Ergonomics Society.

## **11.2 Team Performance and Communication within Networked Human-Machine Systems**

In a previous study we showed that the behavior of individual operators in a networked system involving supervisory control of multiple unmanned air vehicles (UAVs) could be

well characterized and modeled. We explored the utility of both linear (de Visser et al., 2010) and Bayesian probabilistic (Ahmed et al., 2011) models based on the tasking load imposed on the operator—e.g., the number of enemy targets to be handled in an air defense situation, the amount and quality of network message traffic, and individual differences in working memory capacity. However, a key feature of such supervisory control human-machine systems is that human operators typically work in teams, not in isolation. In air defense operations, different operators may be assigned to different monitoring territories and have different areas of responsibility but need to coordinate their actions with one another. One operator within a team can frequently experience a rapid increase in workload as a result of an enemy incursion into his or her area of responsibility. While that operator may require assistance, the immediate demands and stress associated with this rapid increase in workload might hinder that operator's ability to effectively communicate the offloading of tasks to other members within the team. There may also be a cost to individual operators for working with team members in supervisory control tasks. In addition to the cognitive demands placed upon individual operators within a team, increased coordination and communication between team members may be another source of cognitive demand.

Accordingly, this study examined the effects of task load and the reliability of an automated decision aid's message traffic on team performance in a multi-UV simulation of an air defense task (McKendrick et al., 2011). Teams of two operators either received messages that were highly relevant (reliable) to the task they were currently performing, messages that were both relevant and irrelevant (unreliable), or no messages. Team performance was examined under conditions of low and high task load (number of enemy targets to be engaged), as in the previous single-operator study of de Visser et al. (2010). Our measure of team communication focused on the total amount of information conveyed from one teammate to another. We hypothesized that teams would communicate less during high task load. We also envisaged that teams would have higher communication scores when no network messages were provided, but less so when given reliable messages. We predicted that increased scores in "communication detail" would be associated with improved human-system performance. Finally, given that the previous single-operator study of de Visser et al. (2010) found that working memory capacity was a significant contributor to variability in human-system performance, we also obtained verbal and spatial working memory span scores in the present two-person team study, with the expectation that total human-system team performance would also be linked to individual working memory capacity.

Performance was degraded by high task load and improved with an automated decision aid. In addition, team working memory, defined as the average of individual working memory capacity scores, was associated with superior team performance. Higher levels of task load increased the amount of information communicated by teams whereas the presence of an automated decision aid decreased the amount of information communicated by teams. The results have implications for models of team cognition for teams performing similar tasks in a shared, networked human-machine system.

McKendrick, R., Shaw, T., Saqer, H., de Visser, E., & Parasuraman, R. (2011). Team performance and communication within networked supervisory control human-machine systems. In *Proceedings of the Annual Conference of the Human Factors and Ergonomics Society*, Santa Monica, CA.

### **11.3. Adaptive Automation to Improve Human Performance in Supervision of Multiple Uninhabited Aerial Vehicles: Individual Markers of Performance**

Adaptive automation has been shown to offer flexible, context-dependent, and user-specific automation that can enhance human-system performance. While several invocation methods for adaptive automation have been proposed and tested in experimental settings, it is not clear which of these methods can practically be implemented in operational environments. It is therefore important to explore measures that are both predictive of individual performance and that can be easily administered. This study examined both baseline manual performance and working memory capacity to predict future performance with automation (Saqer et al., 2011). Participants were assisted by context-dependent adaptive automation during a simulated command and control task. Results showed that baseline performance without automation predicted overall human-automation performance. Working memory capacity did not predict overall performance, but did predict effective use of the automated aids, so that participants with higher working memory scores used the aids more effectively. These results suggest that effectiveness of human-automation teams can be predicted with quick, cost-efficient, easily measureable markers of performance and can therefore provide practical invocation strategies for adaptive automation.

Saqer, H., de Visser, E., Emfield, A., Shaw, T., & Parasuraman, R. (2011). Adaptive automation to improve human performance in supervision of multiple uninhabited aerial vehicles: Individual markers of performance. In *Proceedings of the Annual Conference of the Human Factors and Ergonomics Society*, Santa Monica, CA.

### **11.4 Measuring Workload using Cerebral Blood Flow during Supervision of Multiple UAVs**

While automated systems have been shown to improve safety and efficiency in operational environments, automation failures can lead to abrupt shifts in workload. Subjective workload scales have been shown to be sensitive to differences in workload, but they are limited in their ability to assess dynamic, moment-to-moment workload variations. Physiological measures may be better suited to assess dynamic workload in complex environments. Such measures can be used to drive adaptive automation. This study explored the feasibility of a relatively new physiological index, Transcranial Doppler Sonography (TCD) as a candidate for adaptive automation studies. Participants performed a long duration task involving supervisory control of multiple UAVs under varying levels of task load. In one group, enemy threats increased once late in the simulation, and in another group enemy threats increased at two points; once early and once late within the simulation. All participants completed a comparison condition in which there was no variation in the number of incoming enemy threats. Cerebral blood

flow velocity (CBFV), as measured by TCD, was measured during task performance. Performance was assessed by the ability of the operator to protect a no-fly zone from enemy incursion. Subjective mental workload was assessed using the NASA-TLX. As performance decreased during periods of high task load, CBFV increased, and there was a close parallel between the CBFV and performance measures. The NASA-TLX was sensitive in detecting differences in workload between the two conditions, but the patterns of results of this subjective measure were insensitive to specific task elements. The results are interpreted in terms of a resource theory of task performance and show that the CBFV measure is sensitive to dynamic changes in task load in complex environments. The findings indicate that CBFV can be used for neuroadaptive automation to support operators supervising multiple UAVs.

Parasuraman, R. (2011). Neuroergonomics: Brain, cognition, and performance at work. *Current Directions in Psychological Science*, 20, 181-186

### **11.5 Effects of Message Modality on Decision Making Performance under Time Pressure**

This study examined the effects of variation in message modality (radio communications vs. text) on decision making performance in a simulated Command and Control Dynamic Targeting Cell. The simulation environment for the experiment was provided by the Distributed Dynamic Decision-Making (DDD) Simulator, version 4. The DDD is a distributed client server simulation that provides a flexible framework in which to study individual and team performance. In general, DDD simulations involve individual (and team) decision-making about complex situations based on information and resources provided by the simulation and other team members. The simulation enables the manipulation of variables such as organizational structure and mission scenario tasking. In addition, a variety of performance measures can be recorded including items such as tasks processed, latencies, and accuracies.

In this study we used a scenario involving a multi-sector air defense environment. Using the appropriate asset for the particular enemy target, participants were tasked with protecting their assigned quadrant of no fly zones by destroying enemy targets that entered it. Once an asset attacked a target it had to be returned to base as an asset; it was not permitted to attack multiple targets. We examined operator performance and workload for participants deploying assets to attack enemy targets and their ability to concurrently monitor auditory or visual communications in three conditions of time pressure (low, medium, and high). Results showed a significant impact in high time-pressure conditions, especially when operators had to process multiple sources of information from the same modality. These findings are a critical step as to understanding multi-tasking performance in command and control environments in general and with regard to communication and spatial monitoring tasks in particular. In collaboration with Cornell University (see below) we plan to model the performance data from this study with a view to developing a basis for adaptive automation to improve human-system performance.

## **12. Consensus Based Real Time Distributed Planning (Lead: MIT)**

Teams of heterogeneous networked agents are regularly employed in autonomous missions (e.g. intelligence, surveillance and reconnaissance (ISR) operations). Typically agents within the team have different roles and responsibilities, and ensuring proper coordination between them is critical for efficient mission execution. However, as the number of agents, system components, and mission tasks increases, planning for such teams becomes increasingly complex, motivating the development of autonomous task allocation and planning methods that improve mission performance. Planning for such teams involves solving complex combinatorial decision problems (NP-Hard), which scale poorly and for which optimal solutions are computationally intractable. The underlying system models typically consist of stochastic, non-linear and time-varying dynamics and constraints, and the planning problem is further complicated by realistic mission considerations such as resource limitations (fuel, payload, bandwidth, etc), asynchronous communication environments, varying network connectivity constraints, and unknown dynamic environments with limited prior information. In this research we address this complex issue of planning for large heterogeneous networked teams by developing computationally efficient robust planning strategies that can effectively account for several of these realistic considerations, such as complex agent models, asynchronous and dynamic communication, and robustness to parameter uncertainty in score functions, transition dynamics, and constraints.

In order to solve realistic planning problems for large heterogeneous networked teams in real time, it is necessary to employ planning algorithms that are computationally efficient and scalable to increasing numbers of agents and tasks. Optimal solution methods for distributing tasks amongst a team of agents are computationally intractable even for moderate sized problems. Many approximation techniques have been considered instead, however, most of these approaches involve centralized planning, which is typically high bandwidth, resource intensive, and slow to react to rapidly changing information. Distributed approaches present several advantages over centralized solutions such as parallelized computation and faster reaction to dynamic environments, however, these often rely on performing consensus on situational awareness among all agents, a process that is often slow and not guaranteed to converge if information about the environment is dynamic.

In this project we developed a real-time distributed planning algorithm, called the Consensus-Based Bundle Algorithm (CBBA), which performs plan consensus in the task space (rather than on situational awareness), providing provably good approximate solutions (both in terms of convergence time and quality) for multi-agent multi-task allocation problems over dynamic networks of heterogeneous agents

### **12.1. Stochastic CBBA Framework**

The CBBA algorithm consists of iterations between two phases, bundle building and consensus. To embed uncertainty models into the CBBA framework, the bundle building phase of CBBA was modified to account for stochasticity in the score functions. In

particular, this involved each agent independently computing bids for tasks using a stochastic score function, and sharing these bids with other agents to determine winning agents and resolve conflicting assignments. The first stochastic metric considered was the expected-value metric, where agents computed the expected value of their score given their assigned task set. In particular, the sequential greedy process to determine which task to add to the current assignment involved each agent computing the marginal contribution to the expected-value score as a result of adding the new task to the current assignment. In missions where stronger performance guarantees than average performance were required, a stochastic metric to mitigate the worst-case possible mission outcomes could be used instead. Here agents could compute marginal scores and compute bids for tasks that maximized their performance in the worst-case. The process of computing the expected-value score or worst-case score of an assignment was nontrivial due to the complex coupling between tasks in an agent's path (for example, taking longer than expected on early tasks impacts the arrival times of subsequent tasks later in the path, thus affecting their scores), and numerical methods were employed to determine how uncertainty would propagate through the agent's assignment execution.

N. Kopeikin, S. S. Ponda, L. B. Johnson, and J. P. How, "Dynamic mission planning for communication control in multiple unmanned aircraft teams," *Unmanned Systems*, vol. 01, no. 01, pp. 41–58, 2013.

## **12.2 Chance-Constrained CBBA**

The previous section described how CBBA was extended to account for stochastic environments by optimizing expected value plans and maximizing worst-case mission performance. In some scenarios however, mitigating worst-case performance is too conservative, and some level of risk may be allowed. An alternate stochastic metric is the chance-constrained metric which provides more flexibility over the conservatism of the solution, while still guaranteeing that the mission performance will be at least as good as the proposed plan value within a certain allowable risk threshold. An issue with this chance-constrained metric however, is that agent scores are coupled through a probabilistic mission constraint and can no longer be optimized individually, limiting its use in distributed planning environments. In this work, we proposed an approximation to the chance-constrained optimization that allowed the problem to be decomposed into distributable chance-constrained sub-problems that could be leveraged within the robust CBBA planning framework. A primary component of this distributed approximation involved allocating individual agent risks given the global mission risk within a consistent framework. Due to the complex coupling between the risk allocation process and the planner assignment selection process, heuristic approximation methods were employed to approximate planner performance given different agent risk allocations. In particular, this work invoked the Central Limit Theorem to employ risk allocation strategies based on Gaussian distributions for both homogeneous and heterogeneous teams. The distributed chance-constrained CBBA algorithm was validated through simulation trials, and results showed large improvements over baseline (deterministic) CBBA, expected-value CBBA, and over worst-case conservative planning strategies, leading to higher mission performance within allowable risk thresholds. Furthermore, the distributed chance-constrained CBBA algorithm achieved similar results to those

obtained by centralized chance-constrained methods, validating the distributed approximation.

S. Ponda, J. Redding, H.L. Choi, J.P. How, M. Vavrina, J. Vian, "Decentralized Planning for Complex Missions with Dynamic Communication Constraints", American Control Conference, 2010, Baltimore, MD

S. Ponda, H.L. Choi, J.P. How, "Predictive Planning for Heterogeneous Human-Robot Teams", AIAA Infotech@Aerospace, 2010, Atlanta, GA

### **12.3. Information-Rich Planning to Reduce Uncertainty**

The previous sections described methods for embedding distribution models into the planning framework to enable robust planning strategies given uncertainty in the environment. A more active approach to handling uncertainty is to use information-based planning strategies to reduce the uncertainty in the environment. The basic notion is that, by actively controlling the measurement process (e.g. sensor locations, vehicle trajectories), model uncertainty can be further reduced through the collection of higher quality data that maximizes information content. We extended the distributed CBBA framework to enable information-rich task allocation through the use of an information-based task heuristic. The approach explicitly considered uncertainty reduction, by computing the Fisher Information associated with different vehicle trajectories and sensing locations, and selecting the task allocations that maximized information content. In joint work with Cornell University, we developed an algorithmic approach to integrate this distributed information-rich CBBA with an information-rich path planning algorithm and with the Cornell information fusion algorithms within a unified architecture with the objective of reducing uncertainty in the target search and tracking process, while considering the complex constraints associated with realistic human-robot search and track missions. In this novel approach, the goal of maximizing information was a primary objective for each of the algorithms at every step, producing a cohesive framework that enabled intelligent and efficient cooperative search and track strategies that were balanced alongside other mission objectives. The resulting task allocation and trajectory planning algorithms were distributed, making the system scalable to large teams of operators and autonomous agents with diverse potential task sets. Furthermore, the information fusion algorithms provided strategies to directly include "soft" inputs from human agents, which were combined with conventional autonomous sensor information via robust particle filtering algorithms, enabling convenient recursive Bayesian updates for efficient replanning. This unified task allocation, trajectory planning and information fusion framework was validated through a set of real-time experiments at Cornell University, involving a human-robot team performing a multi-target search mission, demonstrating the viability of the approach

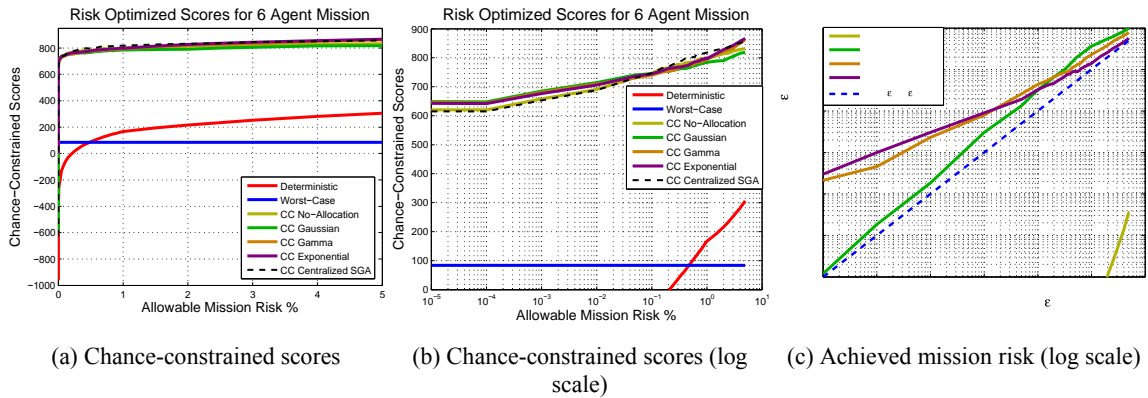
S. S. Ponda, L. B. Johnson, A. Geramifard, and J. P. How, Handbook of Unmanned Aerial Vehicles, ch. Cooperative Mission Planning for Multi-UAV Teams. Springer, 2013.

## 12.4 Risk Allocation Strategies for Distributed Chance-Constrained Task Allocation

The main objective of this project is to address the problem of real-time robust distributed planning for multi-agent networked teams operating in uncertain and dynamic environments. An important issue associated with autonomous planning is that many of the algorithms rely on underlying system models and parameters, which are often subject to uncertainty. This uncertainty can result from many sources including: inaccurate modeling due to simplifications, assumptions, and / or parameter errors; fundamentally nondeterministic processes (e.g., sensor readings, stochastic dynamics); and dynamic local information changes. As discrepancies between the planner models and the actual system dynamics increase, mission performance typically degrades. The impact of these discrepancies on the overall quality of the plan is usually hard to quantify in advance due to nonlinear effects, coupling between tasks and agents, and interdependencies between system constraints (for example, if some tasks take longer than expected this can impact the arrival times of subsequent tasks). However, if uncertainty models of planning parameters are available, they can be leveraged to create robust plans that explicitly hedge against the inherent uncertainty given allowable risk thresholds. This research developed robust distributed task allocation strategies that can be used to plan for multi-agent networked teams operating in stochastic and dynamic environments. In particular, the contributions of this work include: proposing risk allocation strategies that exploit domain knowledge of agent score distributions to improve team performance, providing insights about what stochastic parameters affect the allocations and the overall mission score/performance, and providing results showing improved performance over previously published heuristic techniques in environments with given allowable risk thresholds.

We investigated numerous options for the score function, but of interest in this work is the chance-constrained stochastic metric, which provides probabilistic guarantees on achievable mission performance given allowable risk thresholds and is useful when low probability of mission failure is required.

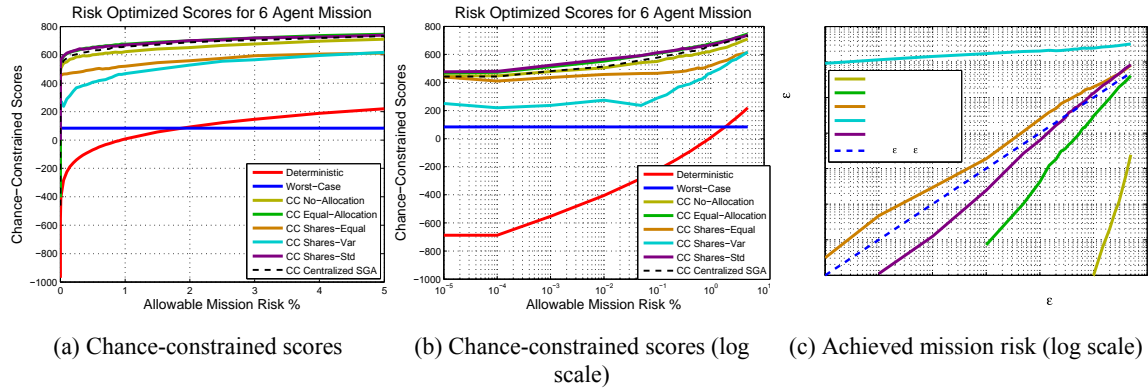
The distributed chance-constrained CBBA algorithm was implemented in simulation for time-critical UAV target tracking missions to validate the risk allocation algorithms.



**Figure 18: Monte Carlo results for a stochastic mission with 6 homogeneous agents and 60 tasks.**



Typical results are shown in Figure 8, which displays Monte Carlo simulation results comparing chance-constrained mission performance for a homogeneous team. The following 7 planning algorithms were compared: a deterministic algorithm (using mean values of parameters), an algorithm optimizing worst-case performance, the chance-constrained CBBA algorithm without explicit risk allocation (all agents planned with mission risk,  $\varepsilon_i = \varepsilon$ , which is typically conservative), chance-constrained CBBA using the different homogeneous risk allocation strategies (Gaussian, Exponential and Gamma), and a centralized chance-constrained sequential greedy algorithm (SGA). The chance-constrained mission scores as a function of mission risk are shown on a linear scale (Figure 8(a)) and a log scale (Figure 8(b)) to highlight performance at low risk levels. The 3 risk allocation strategies achieved higher performance than without risk allocation, with Exponential risk performing best on average. At low risk levels, Gaussian risk gave good performance but as the risk level increased the approximation became worse. All chance-constrained planning approaches performed significantly better than deterministic and worst-case planning which did not account for risk. Figure 8(c) shows the achieved team risk corresponding to the given agent risk allocations  $\varepsilon_i$ , where the dotted line represents a perfect match between desired and actual mission risk. Without risk allocation the team performs conservatively, achieving much lower mission risk than allowed, thus sacrificing performance. With the risk allocation methods, the team is able to more accurately predict the mission risk, where closer matches led to higher scores. Finally, chance-constrained CBBA achieved performance on par with the centralized sequential greedy approach, validating the distributed approximation.



**Figure 19: Monte Carlo results for a stochastic mission with 6 heterogeneous agents and 60 tasks.**


Figure 19 shows results for a heterogeneous stochastic mission where the following 8 planning algorithms were compared: deterministic, worst-case, chance-constrained CBBA without risk allocation, chance-constrained CBBA using an initial risk allocation heuristic proposed with  $H = (2/N_d)^{1/2}$ , chance-constrained CBBA using the heterogeneous Gaussian risk allocation strategies (equal shares, shares based on variance, shares based on std. dev.), and the centralized SGA algorithm. All chance-constrained planning approaches did better than the deterministic and worst-case algorithms. The heterogeneous risk allocation strategy proposed in this paper, with shares proportional to std. dev., performed best overall. Our initial heuristic risk allocation achieved similar performance as well. The other risk allocation approaches performed rather poorly, even though in the equal share case the achieved team risk matched the desired risk well

(Figure (c)). The intuition behind these results is that when agent risk allocations were severely unequal, some agents developed very aggressive plans whereas others selected plans that were too conservative, without considering the effect on the mission as a whole. As a result, the achieved score distributions were quite different between agents, and the convolved mission score distribution yielded lower chance-constrained scores. In general, having a more equitable risk distribution for the team led to higher performing plans. Once again, the performance of CBBA was on par with the centralized approach, validating the distributed approximation.

S. S. Ponda, L. B. Johnson, and J. P. How, “Distributed chance-constrained task allocation for autonomous multi-agent teams,” in American Control Conference (ACC), June 2012.


### 13. Modeling real-time human-automation collaborative scheduling of multiple UVs (Lead: MIT)

A Collaborative Human-Automation Scheduling (CHAS) model was developed using System Dynamics modeling techniques. System Dynamics (SD) is a well-established field that draws inspiration from basic feedback control principles to create simulation models. SD constructs (stocks, flows, causal loops, time delays, feedback interactions) enable investigators to describe and potentially predict complex system performance, which would otherwise be impossible through analytical methods. Through a multi-stage validation process, the CHAS model was tested on three experimental data sets to build confidence in the accuracy and robustness of the model under different conditions.



**Figure 20: Comparison of real-time ratings of trust in the AS (1-7, low to high) throughout the mission between high and low performers. Standard error bars are shown.**

Next, the CHAS model was used to develop recommendations for system design and training changes to improve system performance. These changes were implemented and through an additional set of human subject experiments, the quantitative predictions of the CHAS model were validated. Specifically, test subjects who play computer and video games frequently were found to have a higher propensity to over-trust automation. By priming these gamers to lower their initial trust to a more appropriate level, system



**Figure 21: Predictions using the CHAS model compared to experimental results for gamers.**

performance was improved by 10% as compared to gamers who were primed to have higher trust in the AS. The CHAS model provided accurate quantitative predictions of the impact of priming operator trust on system performance. Finally, the boundary conditions, limitations, and generalizability of the CHAS model for use with other real-time human-automation collaborative scheduling systems were evaluated.

Real-time scheduling in uncertain environments is crucial to a number of domains, especially UV operations. With the ever-increasing demand for UVs for both military and commercial purposes, inverting the operator-to-vehicle ratio will become necessary. Real-time scheduling for multiple UVs in uncertain environments will require the computational ability of optimization algorithms combined with the judgment and adaptability of human supervisors. Despite the potential advantages of human-automation collaboration, inappropriate levels of operator trust, high operator workload, and a lack of goal alignment between the operator and automation can cause lower system performance and costly or deadly errors. The CHAS model can support designers of future UV systems working to address these challenges by simulating the impact of changes in system design and operator training on human and system performance. This could help designers save time and money in the design process, enable the exploration of a wider trade space of system changes than is possible through prototyping or experimentation, and assist in the real-world implementation of multi-vehicle unmanned systems.

A.S. Clare, J.C. Macbeth, and M.L. Cummings, Mixed-Initiative Strategies for Real-time Scheduling of Multiple Unmanned Vehicles, American Control Conference, Montreal, Canada, 2012.

A.S. Clare and M.L. Cummings, Task-Based Interfaces for Decentralized Multiple Unmanned Vehicle Control, Proceedings of AUVSI 2011: Unmanned Systems North America, Washington D.C., August 2011.

### 13.1 Modeling Teamwork of Multi-Human Multi-Agent Teams UVs

A human-in-the-loop experiment was conducted to investigate human-robot agent team structure as well as an agent supporting individual human team members' attention allocation.

USARSim, a robotic simulation performing Urban Search and Rescue (USAR) tasks, was used to provide the underlying simulation for the testbed, as shown in Figure 22. The human operators' tasks were to work as a team of two to explore the unknown environment and identify as many positions of victims as possible.

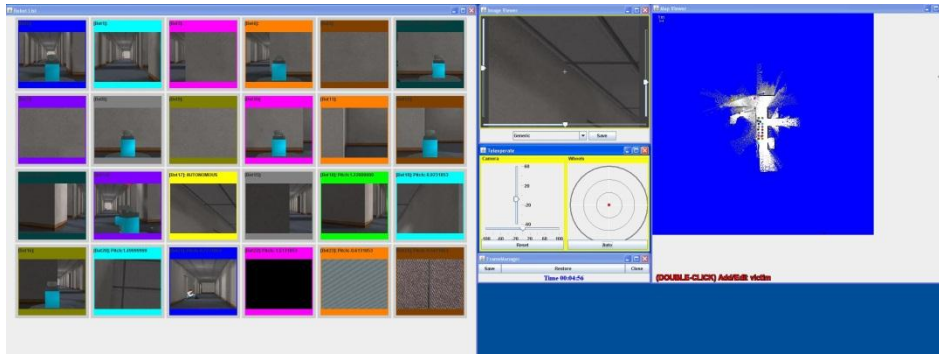


Figure 22: Interface for operating vehicles.

The experiment had two independent variables: *team structure* and *search guidance*. Team structure had two levels:

- *Sector*: each participant controlled 12 robots individually.
- *Shared Pool*: the team shared the control of all 24 robots.

Search Guidance had three levels:

- *Suggested*: system provides a recommendation to switch to another robot when the operator spends thirty seconds on a robot.
- *Enforced*: system provides a recommendation to switch at thirty seconds and switch automatically to another robot five seconds after the recommendation.
- *Off*: system provides no recommendation.

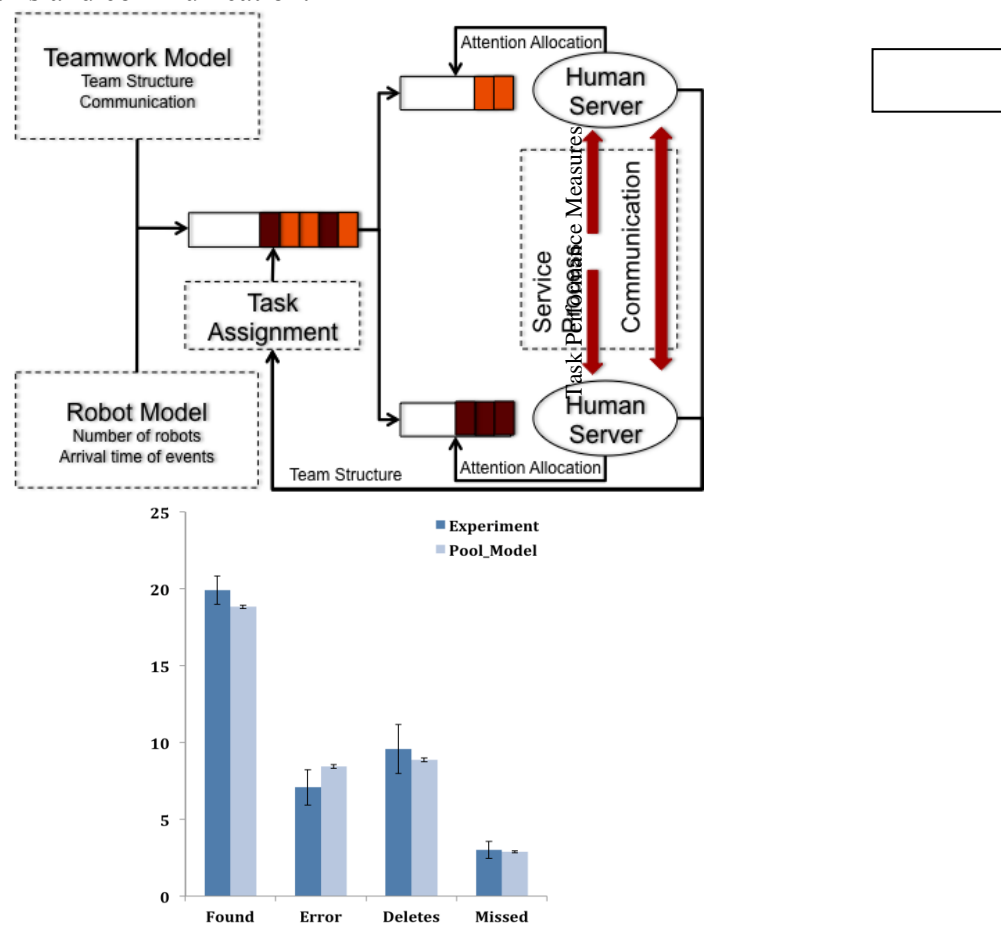
Dependent variables included task performance metrics, subjective workload, operator measures and communication time as team measure. Task performance includes number of victims found, number of errors, number of victims missed and number of deletes.

*Pool* team structure resulted in lower workload ratings than *Sector* team structure, but there was no significant difference in task performance. Further analysis on individual workload and performance suggests that a workload balancing process or back-up behavior occurs in *Pool* teams. *Pool* teams also communicated more while *Sector* teams teleoperated more. Further analyses on communication revealed that communication time was moderately negatively correlated with errors ( $r=-0.309$ ,  $p=0.008$ ) for *Pool* teams,

suggesting that operators in *Pool* teams may correct each other facilitated by communication.

Automated search guidance did not improve or decrease performance, but had an influence on working process. In *Sector* teams, *Suggested* search guidance helped operators mark victims faster when they appeared in the cameras as measured by mean display-to-mark time ( $p=0.024$ ).

A DES model was built based the data and observations from the human-in-the-loop experiment described in the previous section. Operators function as servers in the queuing model and serve the events generated from the robot agents. The overall framework of the model is shown in Figure 23(a).. Four aspects are modeled using DES: arrival process of agent-generated events, service process of human operators, task assignment in teams and communication.



**Figure 23: (a) Discrete Event Simulation Model for Multi-robot Multi-operator. (b) Comparison between DES Model and Experiment for Pool Teams with No Search Guidance.**

Selected publications:

Fei Gao, Andrew S. Clare, Jamie C. Macbeth, M. L. Cummings, “Modeling the Impact of Operator Trust on Performance in Multiple Robot Control,” AAI, 2013.

- Fei Gao, M.L. Cummings, "Using Discrete Event Simulation to Model Multi-Robot Multi-Operator Teamwork," In Proceedings of the Human Factors and Ergonomics Society Annual Meeting, SAGE Publication, Vol. 56, No. 1, pp. 2093-209, Boston, MA, October, 2012.
- Fei Gao, Missy L. Cummings, and Luca F. Bertuccelli, "Teamwork in controlling multiple robots," In Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction (HRI '12), ACM, New York, NY, USA, 81-8, 2012. DOI=10.1145/2157689.2157703

## Personnel Supported

8 faculty, 4 postdocs, 3 research staff, 17 graduate students, 2 undergraduates

## Degrees and degree type granted

Prasanna Velagapudi, PhD in Robotics, August 2012. Thesis: "Distributed Planning Under Uncertainty for Large Teams", Carnegie Mellon University

[http://www.cs.cmu.edu/~softagents/theses/Velagapudi\\_PhD\\_Thesis\\_2012.pdf](http://www.cs.cmu.edu/~softagents/theses/Velagapudi_PhD_Thesis_2012.pdf)

Steven Okamoto, PhD in Computer Science, August 2012. Thesis: "Allocating Virtual and Physical Flows for Multiagent Team in Mutable, Networked Environments", Carnegie Mellon University

[http://www.cs.cmu.edu/~softagents/theses/Okamoto\\_PhD\\_Thesis\\_2012.pdf](http://www.cs.cmu.edu/~softagents/theses/Okamoto_PhD_Thesis_2012.pdf)

Huadong Wang PhD in Information Sciences, June 2013 Thesis "Asynchronous Visualization of Spatiotemporal Information for Multiple Moving Targets", University of Pittsburgh

<http://d-scholarship.pitt.edu/19545/>

Ewart de Visser, PhD, July 2012. Thesis: "the World is not Enough: Trust in Cognitive Agents", George Mason University

[http://www.cs.cmu.edu/~softagents/theses/GMU\\_deVisser\\_PhD\\_Thesis\\_2012.pdf](http://www.cs.cmu.edu/~softagents/theses/GMU_deVisser_PhD_Thesis_2012.pdf)

Nisar Ahmed, PhD in Mechanical Engineering, January 2012, Cornell "Probabilistic Modeling and Estimation with Human Inputs in Semi-Autonomous Systems"

<http://ecommons.library.cornell.edu/handle/1813/29445>

Also at: [http://www.cs.cmu.edu/~softagents/theses/Cornell\\_Ahmed\\_PhD\\_Thesis\\_2012.pdf](http://www.cs.cmu.edu/~softagents/theses/Cornell_Ahmed_PhD_Thesis_2012.pdf)

Danelle Shah, PhD in Mechanical Engineering, January 2012, Cornell "Towards Natural and Robust Human-Robot Interaction Using Sketch and Speech"

<http://ecommons.library.cornell.edu/handle/1813/29272>

Also at: [http://www.cs.cmu.edu/~softagents/theses/Cornell\\_Shah\\_PhD\\_Thesis\\_2012.pdf](http://www.cs.cmu.edu/~softagents/theses/Cornell_Shah_PhD_Thesis_2012.pdf)

Andrew Clare, PhD in Aeronautical Engineering, June 2013. "Modeling real time Human Automation Collaborative Scheduling of Unmanned Vehicles" MIT

[http://web.mit.edu/aeroastro/labs/halab/papers/Clare\\_Doctoral\\_Thesis.pdf](http://web.mit.edu/aeroastro/labs/halab/papers/Clare_Doctoral_Thesis.pdf) .

Sameera Ponda, PhD in Aeronautical Engineering, September 2012 Thesis: "Robust Distributed Planning Strategies for Autonomous Multi-Agent Teams", MIT

[http://acl.mit.edu/papers/Ponda\\_PhD\\_Thesis\\_Final.pdf](http://acl.mit.edu/papers/Ponda_PhD_Thesis_Final.pdf).

Nathan Brooks MS in Robotics August 2011 Thesis:"Scalable Target Detection for Large Robot Teams" Carnegie Mellon University.

[http://www.cs.cmu.edu/~softagents/theses/Brooks\\_MS\\_thesis\\_2011.pdf](http://www.cs.cmu.edu/~softagents/theses/Brooks_MS_thesis_2011.pdf)

Bruno Hexsel, MS in Robotics, December 2010 Thesis “Coverage Control for Mobile Anisotropic Sensor Networks” CMU

[http://www.cs.cmu.edu/~softagents/theses/Hexsel\\_MS\\_thesis\\_2010.pdf](http://www.cs.cmu.edu/~softagents/theses/Hexsel_MS_thesis_2010.pdf)

Siddharth Mehrotra MS in Robotics December 2010 Thesis “Effects of Robot Self-Reflection on Operator Control of Robot Teams”, CMU

[http://www.cs.cmu.edu/~softagents/theses/Mehrotra\\_MS\\_thesis\\_2010.pdf](http://www.cs.cmu.edu/~softagents/theses/Mehrotra_MS_thesis_2010.pdf)

Breelyn Kane MS in Robotics January 2010 Thesis “The Operator, a Valuable Resource: Asking for Help Through Adaptive Autonomy”, CMU

[http://www.cs.cmu.edu/~softagents/theses/Kane\\_MS\\_thesis\\_2010.pdf](http://www.cs.cmu.edu/~softagents/theses/Kane_MS_thesis_2010.pdf)

Lingzhi Luo, MS in Robotics May 2012 Thesis “Distributed Algorithms for Constrained Multi-Robot Dynamic Task Assignment with Formal Performance Guarantees”, CMU

[http://www.cs.cmu.edu/~softagents/theses/Luo\\_MS\\_Thesis\\_2012.pdf](http://www.cs.cmu.edu/~softagents/theses/Luo_MS_Thesis_2012.pdf)

## Honors/Awards

Philip Walker, Steven Nunnally, Nilanjan Chakraborty, Michael Lewis, Katia Sycara, “Levels of Automation for Human Influence of Robot Swarms”, HFES, , San Diego, September 30-October 4 2013. (*Best Paper Award*)

Katia Sycara: Co-recipient (second time in a row) of the Semantic Web Scientific Association *most influential 10-year paper award*, for the paper titled "Semantic Matching of Web Services Capabilities.". The award was presented at the 11<sup>th</sup> International Semantic Web Conference (ISWC), Boston, USA, November 11-15, 2012.

Katia Sycara: Co-recipient of the Semantic Web Scientific Association *most influential 10-year paper award*, for the paper titled "DAML-S: Semantic Markup for Web Services". The award got presented at the 10<sup>th</sup> International Semantic Web Conference (ISWC), Bonn, Germany, October 23-27, 2011.

Ryan McKendrick (October 2011): *Best Student Paper Award* for paper “Team performance and communication within networked supervisory control human-machine systems.” Cognitive Engineering and Decision Making Technical Group. Human Factors and Ergonomics Society

Raja Parasuraman (October 2011): *Pioneer in Human-Automation Research Award*, Cognitive Engineering and Decision Making Technical Group, Human Factors and Ergonomics Society.

Raja Parasuraman (February 2012): *Outstanding Educator Award*, International Ergonomics Association Triennial Award

*Best paper award* at the AIAA Guidance, Navigation and Control Conference, 2011.



Ponda, S., Ahmed, N., Luders, B., Sample, E., Levine, D., Hoossainy, T., Shah, D., Campbell, M., and How, J. P., “Decentralized Information-Rich Planning and Hybrid Sensor Fusion for Uncertainty Reduction in Human-Robot Missions,” *AIAA Guidance, Navigation and Control Conference*, 2011.

*Best paper award*, conference on Behavior Representation in Modeling and Simulation (BRIMS 2009) David Reitter and Christian Lebiere. A subsymbolic and visual model of spatial path planning. In: *Proc. Behavior Representation in Modeling and Simulation (BRIMS)*, 2009.

ACT-UP model of dynamic control was the *winner*, of Dynamic Stocks and flows cognitive modeling competition (2010). Reitter, D. Metacognition and multiple strategies in a cognitive model of online control. *Journal of Artificial General Intelligence*, 2(2):20-37, 2010.

Katia Sycara, selected as member of the National Academies Study “From Data to Decision: Integrating Human, Machines and Networks”, 2012-2014.

Missy Cummings selected as member of the National Academies Study “From Data to Decision: Integrating Human, Machines and Networks”, 2012-2014.

Mark Campbell was selected to the Defense Sciences Study Group (DSSG), 2012-2013.

## **Publications**

### **Journals**

1. Ahmed, N., de Visser, E., Shaw, T., Mohamed-Ameen, A., Campbell, M. A., & Parasuraman, R. Predicting human-automation performance in networked systems using statistical models: The role of working memory capacity. *Interacting with Computers* (in press).
2. W. Whitacre, M. Campbell, “Information-Theoretic Optimization of Periodic Orbits for Persistent Cooperative Geolocation,” accepted Nov 2011, and to appear in the *AIAA Journal of Guidance, Control, and Dynamics*.
3. McKendrick, R., Shaw, T., de Visser, E., Saqer, H., Kidwell, B., & Parasuraman, R. (accepted). Team performance in networked supervisory control of unmanned air vehicles: Effects of automation, working memory and communication content. *Human Factors*, doi:10.1177/0018720813496269.
4. Parasuraman, R., Kidwell, B., Olmstead, R., Lin, M-K., Jankord, R., & Greenwood, P. (accepted). Interactive effects of the COMT gene and training on individual differences in supervisory control of unmanned vehicles. *Human Factors*, in press.
5. Kolling, A., Sycara, K., Nunnally, S., Lewis, M. Human Swarm Interaction: An Experimental Study of Two Types of Interaction with Foraging Swarms, *Journal of Human-Robot Interaction*, June 2013.

6. Kleiner, A., Kolling, A., Lewis, M., & Sycara, K. (2013) Hierarchical Visibility for Guaranteed Search in Large-Scale Outdoor Terrain, *Journal of Autonomous Agents and Multi-Agent Systems*, 26(1), 1-36.
7. N. Kopeikin, S. S. Ponda, L. B. Johnson, and J. P. How, "Dynamic mission planning for communication control in multiple unmanned aircraft teams," *Unmanned Systems*, vol. 01, no. 01, pp. 41–58, 2013.
8. Hancock, P. A., Jagacinski, R., Parasuraman, R., Wickens, C., Wilson, G., & Kaber, D. (2013). Human-automation interaction research: Past, present and future. *Ergonomics in Design*, 21(2), 9-14.
9. Whitacre, William W., and Mark E. Campbell. "Cooperative Estimation Using Mobile Sensor Nodes in the Presence of Communication Loss." *Journal of Aerospace Information Systems* 10.3 (2013): 114-130.
10. N. Ahmed, M. Campbell, "On Estimating Simple Probabilistic Discriminative Models with Subclasses," *Expert Systems With Applications*, Vol 39, No 7, June 2012, pp 6659–6664.
11. S. S. Ponda, L. B. Johnson, A. N. Kopeikin, H.-L. Choi, and J. P. How, "Distributed planning strategies to ensure network connectivity for dynamic heterogeneous teams," *IEEE Journal on Selected Areas in Communications*, vol. 30, pp. 861 – 869, June 2012.
12. A.S. Clare, P.C.P. Maere, and M.L. Cummings, Assessing Operator Strategies for Real-time Replanning of Multiple Unmanned Vehicles, *Intelligent Decision Technologies*, Vol. 6, No. 3, pp. 221-231, 2012.
13. De Visser, E., & Parasuraman, R. (2011). Adaptive aiding of human-robot teaming: Effects of imperfect automation on performance, trust, and workload. *Journal of Cognitive Engineering and Decision Making*, 5, 209-231.
14. Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y.C., de Visser, E., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors*, 53, 517-727.
15. Parasuraman, R. (2011). Neuroergonomics: Brain, cognition, and performance at work. *Current Directions in Psychological Science*, 20, 181-186
16. Lewis, M., Wang, H., Chien, S., Valagapudi, P., Scerri, P. Sycara, K Process and Performance in Human-Robot Teams, *Journal of Cognitive Engineering and Decision Making*, 5(2), 186-208, 2011
17. David Reitter, Frank Keller, and Johanna D. Moore. A computational cognitive model of syntactic priming. *Cognitive Science*, 35(4):587-637, 2011.
18. David Reitter and Christian Lebiere. How groups develop a specialized domain vocabulary: A cognitive multi-agent model. *Cognitive Systems Research*, 12(2):175-185, 2011.
19. Lebiere, C., & Anderson, J. R. (2011). Cognitive constraints on decision making under uncertainty. *Frontiers in Cognition* 2 (305).
20. D. C. Shah, J. R. Schneider, M. E. Campbell, "Robust, Sketch based Command and Control of Robot Teams," *Proceedings of the IEEE*, published on-line Dec 2011, Vol 100, No. 3, pp 604-622.
21. W. Whitacre, M. Campbell, "Decentralized Geolocation and Sensor Bias Estimation for UAVs with Articulating Cameras," *AIAA Journal of Guidance, Controls, and Dynamics*, Vol 34, No 2, Mar-Apr 2011, pp 564-573.

22. Ahmed, N. and Campbell, M., "Variational Bayesian Learning of Probabilistic Discriminative Models with Latent Softmax Variables," *IEEE Transactions on Signal Processing*, Vol 59, No 7, July 2011, pp 3143-3154.
23. Scerri, P., Ma, Z., Chien, S., Wang, H., Lee, P., Lewis, M. & Sycara, K. (2011) An initial evaluation of approaches to building entry for large robot teams, *Journal of Intelligent and Robotic Systems*, 64, 145-159.
24. Lewis, M., Wang, H., Chien, S., Velagapudi, P., Scerri, P. & Sycara, K. (2010). Choosing autonomy modes for multirobot search, *Human Factors*, 52(2), 225-233.
25. Reitter, D. (2010). Metacognition and multiple strategies in a cognitive model of online control. *Journal of Artificial General Intelligence*, 2(2), 20-37.
26. Gclinton, R., Sycara, K., Scerri, D., & Scerri, P. (2010). The statistical mechanics of belief sharing in multi-agent systems. *Information Fusion*, 11(3), 256-266.
27. Gluck, K. A., Stanley, C. T., Moore, L. R., Reitter, D., & Halbrügge, M. (2010). Exploration for understanding in model comparisons. *Journal of Artificial General Intelligence*, 2(2):88-107, 2010.
28. Reitter, D. & Lebiere, C. (2010). A cognitive model of spatial path planning. *Journal of Computational and Mathematical Organization Theory*, 16, 220-245.
29. Lewis, M., Sycara, K., & Scerri, P. . Scaling up wide-area-search-munition teams, *IEEE Intelligent Systems*, 24(3), 10-13, 2009.
30. Sycara, K., Gclinton R., Yu B., Giampapa, J., Owens S., Lewis, M., Grindle, C., "An Integrated Framework for High Level Information Fusion", *International Journal Information Fusion*. Vol 10, Issue 1, January 2009, pp 25-50.
31. Parasuraman, R., Cosenzo, K., & de Visser, E. (2009). Adaptive automation for human supervision of multiple uninhabited vehicles: Effects on change detection, situation awareness, and mental workload. *Military Psychology*, 21.270-297.
32. Choi, H. Bruent, L. and How, J. "Consensus-based decentralized auctions for robust task allocation," *IEEE Transactions on Robotics*, 2009.
33. Cummings M. and Mitchell, P. "Predicting controller capacity in remote supervision of multiple unmanned vehicles," *IEEE Systems, Man, and Cybernetics, Part A Systems and Humans*, vol. 38, no. 2, pp. 451-460, 2008.
34. Nehme, C Mekdeci, B. Crandall, J. and Cummings, M. "The impact of heterogeneity on operator performance in future unmanned vehicle systems," *The International Command and Control Journal*, vol. 2, 2008.

## **Book Chapters**

35. Lewis, M. "Human Interaction with Multiple Remote Robots," In D. Kaber (Ed), *HF Reviews Volume 9, on Human Performance in Teleoperation and Beyond*, HFES (in press).
36. S. S. Ponda, L. B. Johnson, A. Geramifard, and J. P. How, *Handbook of Unmanned Aerial Vehicles*, ch. Cooperative Mission Planning for Multi-UAV Teams. Springer, 2013 (in press).
37. N. Kopeikin, S. S. Ponda, and J. P. How, *Handbook of Unmanned Aerial Vehicles*, ch. Control of Communication Networks for Teams of UAVs. Springer, 2013 (in press).

38. Scerri, P., Velagapudi, P., Sycara, K. "Analyzing the Theoretical Performance of Information Sharing", In Dynamics of Information Systems: Theory and Applications, Springer 2010.
39. Lewis, M. & Wang, J. (2010). Coordination and automation for controlling robot teams, in M. Barnes & F. Jentsch (Eds), Human-Robot Interactions in Future Military Operations, Burlington, VT: Ashgate, 397-418.
40. Lewis, M. and Wang, J. (2009). Assessing coordination demand in cooperating robots, In Madhavan, Tunstel and Messina (Ed.) Performance Evaluation and Benchmarking of Intelligent Systems, Springer, New York, 169-186.
41. Paruchuri, P., Grinton R., Sycara, K. Scerri, P. "Effect of Humans on Belief propagation in large heterogeneous teams", in Hirsch, M Pardalos, P. and Murphy R (eds), Dynamics of Information Systems, Springer, 2009.
42. Grinton, R., Paruchuri, P., Scerri, P. Sycara, K "Self-organized criticality of belief propagation in large heterogeneous teams", Hirsch, M Pardalos, P. and Murphy R (eds), Dynamics of Information Systems, Springer, 2009.
43. Velagapudi, P., Prokopiiev, O., Scerri, P., Sycara, K. "A Token-Based Approach to Sharing Beliefs in a Large Multiagent Team", Control and Information systems, Springer, 2009.

### **Conference Papers**

44. Luo, L., Chakraborty, N. and Sycara, K. "Distributed Algorithm Design for Multi-robot Generalized Task Assignment Problem", Proceedings of International Conference on Intelligent Robots and Systems (IROS), Tokyo, Japan, November 3-8, 2013.
45. Nunnally, S., Walker, P., Chakraborty, N., Lewis, M. Sycara, K Using Haptic Feedback in Human Robotic Swarm Interaction, HFES, San Diego, September 30-October 4 2013
46. Walker, P, Nunnally, S., Chakraborty, N., Lewis, M, Sycara, K. Levels of Automation for Human Influence of Robot Swarms, HFES, San Diego, September 30-October 4 2013. CSTG Best Paper Award.
47. Nunnally, S., Walker, P., Chakraborty, N., Lewis, M. Sycara, K Using Coverage for Measuring the Effect of Haptic Feedback in Human Robot Swarm Interaction, In Proceedings of the Conference on System Man and Cybernetics, Manchester, UK., October 13-16 2013.
48. Walker, P., Annipur Amraii, S., Lewis, M., Chakraborty, N., Sycara, K. Human Control of Leader-Based Swarms Proceedings of the Conference on System Man and Cybernetics, Manchester, UK., October 13-16 2013.
49. Zadorozhny, V., Lewis, M. (2013). Information Fusion based on Collective Intelligence for Multi-Robot Search and Rescue Missions. To appear in Proceedings of the 14th International Conference on Mobile Data Management (MDM'13)
50. Zadorozhny, V., Lewis, M. (2013). Information Fusion for USAR Operations Based on Crowdsourcing, Proceedings of the 16th International Conference on information Fusion, Istanbul, TK

51. Goerner, J, Chakraborty, N., Sycara, K., Energy Efficient Data Collection with Mobile Robots in Heterogeneous Sensor Networks, International Conference on Robotics and Automation (ICRA), Karlsruhe, Germany, May 6-10, 2013.
52. Luo, L., Chakraborty, N., Sycara, K. Distributed Algorithm Design for Multi-Robot Task Assignments with Deadlines for Tasks, International Conference on Robotics and Automation (ICRA), Karlsruhe, Germany, May 6-10, 2013
53. S. S. Ponda, L. B. Johnson, and J. P. How, "Risk allocation strategies for distributed chance-constrained task allocation," in American Control Conference (ACC), June 2013,
54. Fei Gao, "Modeling teamwork of multi-human multi-agent teams," In Proceedings of the 2013 conference on Computer supported cooperative work companion, pp. 47-50, San Antonio, Texas, February, 2013.
55. Fei Gao, Andrew S. Clare, Jamie C. Macbeth, M. L. Cummings, "Modeling the Impact of Operator Trust on Performance in Multiple Robot Control," AAAI, 2013.
56. Reitter, D., & Scerri, P. (2013). Smooth dynamics, good performance in cognitive-agent congestion problems. In *Proceedings of the 35th Annual Meeting of the Cognitive Science Society*.
57. Reitter, D., & Scerri, P. (2013). Cognitive instance-based learning agents in a multi-agent congestion game. In *Proceedings of Workshop on Information Sharing in Large Scale Multi-Agent Systems, at AAMAS 2013*.
58. Luo, L., Chakraborty, N., Sycara, K., "Competitive Analysis of Repeated Greedy Auction Algorithm for Online Multi-Robot Task Assignment", International Conference on Robotics and Automation (ICRA), St. Paul, Minnesota, May 14-18, 2012.
59. Okamoto, S., Hazon, N., Sycara, K. "Solving Non-Zero Sum Multiagent Network Flow Security Games with Attack Costs", International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-12), Valencia, Spain, June 4-8, 2012.
60. Varakantham, P., Yeoh, W., Velagapudi, P., Sycara, K., Scerri, P. "Prioritized Shaping of Models for Solving DEC-POMDPs" International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-12), Valencia, Spain, June 4-8, 2012
61. Linglong Zhu, Yang Xu, Paul Scerri, Han Liang, "An Information Sharing Algorithm For Large Dynamic Mobile Multi-agent Teams" International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Valencia, Spain, June 4-8, 2012.
62. David Reitter and Christian Lebiere. Social cognition: Memory decay and adaptive information filtering for robust information maintenance. In Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12), 2012.
63. Lewis, M. and Wang, H., Kolling, A., Sycara, K., Brooks, N. & Scerri, P. Asynchronous Displays for multi-UV Search Tasks. In: AAIA InfoTech, 2012.
64. Chien, S., Mehrotra, S., Lewis, M. & Sycara, K. Effects of Unreliable Automation in Scheduling Operator Attention for Multi-Robot Control, 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC'12), Oct 14-17, Seoul, Korea, 2012

65. Nunnally, S., Walker, P., Lewis, M., Kolling, A., Chakraborty, N., Sycara, K. & Goodrich, M. Human Influence of Robotic Swarms with Bandwidth and Localization Issues, 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC'12), Oct 14-17, Seoul, Korea, 2012
66. Walker, P., Kolling, A., Chakraborty, N., Nunnally, S., Sycara, K. & Lewis, M.. Neglect Benevolence in Human Control of Swarms in the Presence of Latency, 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC'12), Oct 14-17, Seoul, Korea, 2012.
67. Chien, S., Mehrotra, S., Brooks, N., Lewis, M. & Sycara, K. Scheduling operator attention for multi-robot control, 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'12), October 7-12, Vilamoura, Portugal, 2012
68. Kolling, A., Nunnally, S., Lewis, M., and Sycara, K., Towards human control of robot swarms, International Conference on Human-Robot Interaction (HRI'12), March 5-8, 2012, Boston, MA., 89-96.
69. Meneguzzi, F., Oh, J., Chakraborty, N., Sycara, K., Mehrotra, S., Tittle, J., Lewis, M. "A cognitive architecture for emergency response", International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Valencia, Spain, June 4-8, 2012.
70. D. Reitter and R. Levy, editors. Proc. 3rd Workshop on Cognitive Modeling and Computational Linguistics. Association for Computational Linguistics, Montréal, Quebec, Canada, 2012.
71. J. Diesner and D. Reitter, editors. Proc. Words and Networks 2012: Language Use in Socio-Technical Networks (WON2012). Association for Computing Machinery, Chicago, IL, 2012.
72. David Reitter. Lexical language evolution in networked human groups. In Words and Networks: Language Use in Socio-Technical Networks (WON 2012), Chicago, IL, 2012.
73. Ying Xu, Tinglong Dai, Katia Sycara, Michael Lewis. 2012. A Mechanism Design Model in Multi-Robot Service Queues with Strategic Operators and Asymmetric Information. Proceedings of the 51st IEEE Conference on Decision and Control: CDC'12.
74. S. Ponda, L. Johnson, J. P. How, Distributed Chance-Constrained Task Allocation for Autonomous Multi-Agent Teams", American Control Conference (ACC), June 2012
75. R. Tse, N. Ahmed, M. Campbell, "Unified Mixture-Model Based Terrain Estimation with Markov Random Fields," 2012 IEEE International Conference on Multisensor Fusion and Integration.
76. N. Ahmed, J. Schoenberg, M. Campbell, "Fast Weighted Exponential Product Rules for Robust General Multi-Robot Data Fusion," Robotics Science and Systems Conference, 2012.
77. E. Sample, N. Ahmed, M. Campbell, "An Experimental Evaluation of Bayesian Soft Human Sensor Fusion in Robotic Systems," 2012 AIAA Guidance, Navigation and Control Conference.

78. L. B. Johnson, H.-L. Choi, S. S. Ponda, and J. P. How, "Allowing non-submodular score functions in distributed task allocation," in IEEE Conference on Decision and Control (CDC), Dec 2012.
79. T. Campbell, S. S. Ponda, G. Chowdhary, and J. P. How, "Planning under uncertainty using nonparametric Bayesian models," in AIAA Guidance, Navigation, and Control Conference (GNC), August 2012.
80. Fei Gao, M.L. Cummings, "Using Discrete Event Simulation to Model Multi-Robot Multi-Operator Teamwork," In Proceedings of the Human Factors and Ergonomics Society Annual Meeting, SAGE Publication, Vol. 56, No. 1, pp. 2093-209, Boston, MA, October, 2012.
81. Fei Gao, Missy L. Cummings, and Luca F. Bertuccelli, "Teamwork in controlling multiple robots," In Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction (HRI '12), ACM, New York, NY, USA, 81-8, 2012. DOI=10.1145/2157689.2157703
82. A.S. Clare, J.C. Macbeth, and M.L. Cummings, Mixed-Initiative Strategies for Real-time Scheduling of Multiple Unmanned Vehicles, American Control Conference, Montreal, Canada, 2012.
83. A.S. Clare, M.L. Cummings, and L. Bertuccelli, Identifying Suitable Algorithms for Human-Computer Collaborative Scheduling of Multiple Unmanned Vehicles, AIAA Aerospace Sciences Meeting, Nashville, TN, 2012.
84. Kidwell, B., Calhoun, G., Ruff, R., & Parasuraman, R. (2012,). Adaptable and adaptive automation for supervisory control of multiple autonomous vehicles. Proceedings of the Annual Conference of the Human Factors and Ergonomics Society, Santa Monica, CA: Human Factors and Ergonomics Society.
85. Satterfield, K., Ramirez, R., Shaw, T., & Parasuraman, R. (2012). Measuring workload during a dynamic supervisory control task using cerebral blood flow velocity and the NASA-TLX. Proceedings of the Annual Conference of the Human Factors and Ergonomics Society, Santa Monica, CA: Human Factors and Ergonomics Society.
86. Steven Okamoto\_, Praveen Paruchuri, Yonghong Wang, Katia Sycara, Janusz Marecki and Mudhakar Srivatsa "Multiagent Communication Security in Adversarial Settings", International Conference on Intelligent Agent Technology, Lyon, France, August 22-27, 2011.
87. Yonghong Wang, Katia Sycara and Paul Scerri Towards an Understanding of the Value of Cooperation in uncertain world, International Conference on Intelligent Agent Technology, Lyon, France, August 22-27, 2011.
88. Yonghong Wang, Katia Sycara and Paul Scerri Multi-Variate Distributed Data Fusion with Expensive Sensor Data. International Conference on Intelligent Agent Technology, Lyon, France, August 22-27, 2011.
89. Lee, P., Kolling, A. and Lewis, M. Combining latency and utilization in investigating human operator workload, 2011 IEEE International Conference on Systems, Man, and Cybernetics, (SMC'11) October, 2011, Anchorage, Alaska, USA.
90. Abedin, S., Wang, H., Lee, P., Lewis, M., Brooks, N., Owens, S., Scerri, P. and Sycara, K. SUAVE: Integrating UAV Video Using a 3D Model 2011 IEEE International Conference on Systems, Man, and Cybernetics, (SMC'11) October, 2011, Anchorage, Alaska, USA.

91. Tinglong Dai, T., Sycara, K. & Lewis, M. (2011) A game theoretic queuing approach to self-reflection in decentralized human-robot interaction systems, 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'11), September 25-30, 2011, San Francisco, CA.
92. Kolling, A., Kleiner, A., Lewis, M. & Sycara, K. (2011) Computing and executing strategies for moving target search, 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'11), September 25-30, 2011, San Francisco, CA.
93. Chien, S, Wang, H. Lewis, M., Mehrotra, S., Sycara, K. Effects of Alarms on Control of Robot Teams, Proceedings of the Human Factors and Ergonomic Society 55th Annual Meeting (HFES-2011), September 19-23, Las Vegas, NV pp.434-438.
94. Abedin, S., Brooks, N., Owens, S., Scerri, P., Lewis, M., & Sycara, K. SUAVE: Integrating UAV Video Using a 3D Model, Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11 Las Vegas, NV).
95. Brooks, N., Wang, H., Chien, S., Lewis, M., Scerri, P., & Sycara, K. Asynchronous Control with ATR for Large Robot Teams, Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11), September 19-23, Las Vegas, NV, 2011..
96. Wang, H., Kolling, A., Brooks, N., Lewis, M. & Sycara, K. (2011). Synchronous vs. Asynchronous Control for Large Robot Teams. Proceedings of Human Computer International Part II (14) (HCI'11), 415-424.
97. David Reitter and Christian Lebiere. Towards cognitive models of communication and group intelligence. In Proceedings of the 33rd Annual Meeting of the Cognitive Science Society, pages 734-739, Boston, MA, July 2011.
98. David Reitter, Katia Sycara, Christian Lebiere, Yury Vinokurov, Antonio Juarez, and Michael Lewis. How teams benefit from communication policies: information flow in human peer-to-peer networks. In Proceedings of the 20th Behavior Representation in Modeling & Simulation (BRIMS), 2011.
99. W. Mason, D. Reitter, A. Coman, and B. Hirst. Cognition and social dynamics: a new approach to emergent phenomena. Symposium at the 23rd Annual Convention of the Association for Psychological Science, Washington, D.C., 2011.
100. Frank Keller and David Reitter, editors. Proc. 2nd Workshop on Cognitive Modeling and Computational Linguistics. Association for Computational Linguistics, Portland, OR, USA, 2011.
101. McKendrick, R., Shaw, T., Saqer, H., de Visser, E., & Parasuraman, R. (2011). Team performance and communication within networked supervisory control human-machine systems. In Proceedings of the Annual Conference of the Human Factors and Ergonomics Society, Santa Monica, CA.
102. Wang, H., Kolling A., Abedin, S., Lee, P., Chien, S., Lewis M, Books, N., Owens, S., Scerri, P., Sycara, K Scalable target detection for large robot teams Proceedings of the 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI-2011), Lausanne, Switzerland, March 3-9, 2011.
103. Lewis, M, Sycara, K. Effects of automation on situation awareness in controlling robot teams, The Fourth International Conference in Advances in Human Computer Interaction, Guadeloupe, February 23-27, 2011.



104. Lewis, M. & Sycara, K.. Network centric control for multirobot teams in urban search and rescue, Proceedings of the 2011 Hawaii International Conference on Systems Sciences (HICSS-44), January 4-7, 2011.
105. Dai, T, Sycara, K., Lewis, M. A game theoretic queuing approach to self-assessment in human-robot interaction systems. IEEE International Conference on Robotics and Automation (ICRA 2011), May 9-13, Shanghai, China, 2011.
106. Kolling, A., Kleiner, A., Lewis, M. Sycara, K. Computing and executing strategies for multi-robot search. IEEE International Conference on Robotics and Automation (ICRA 2011), May 9-13, Shanghai, China, 2011.
107. Luo, L., Chakraborty, N., Sycara, K., Multi-robot algorithm for tasks with set precedence constraints, IEEE International Conference on Robotics and Automation (ICRA 2011), May 9-13, Shanghai, China, 2011.
108. Hexsel B., Chakraborty, N., Sycara, K. Coverage control for mobile anisotropic sensor networks, IEEE International Conference on Robotics and Automation (ICRA 2011), May 9-13, Shanghai, China, 2011.
109. Velagapudi, P., Varakantham P., Scerri, P., Sycara K., Distributed Model Shaping for Scaling to Decentralized POMDPs with hundreds of agents, Proceedings of the 2011 Conference on Autonomous Agents and Multi-Agent Systems, May 2-6, Taipei, Taiwan, 2011.
110. Hofmann, L., Chakraborty, N., Sycara, K., The Evolution of Cooperation in Self-Interested Agent Societies: A Critical Study, Proceedings of the 2011 Conference on Autonomous Agents and Multi-Agent Systems, May 2-6, Taipei, Taiwan, 2011.
111. Glington, R., Scerri, P., Sycara K., An Investigation of the Vulnerabilities of Scale Invariant Dynamics in Large Teams Proceedings of the 2011 Conference on Autonomous Agents and Multi-Agent Systems, May 2-6, Taipei, Taiwan, 2011.
112. Zivan, R., Paruchuri, P., Sycara, K., Dudik M., Maximizing revenue in symmetric resource allocation when user utilities exhibit diminishing returns (short paper) Proceedings of the 2011 Conference on Autonomous Agents and Multi-Agent Systems, May 2-6, Taipei, Taiwan, 2011.
113. Okamoto, S., Brooks, N., Owens, S., Scerri, P., Sycara, K., Allocating Spatially Distributed Tasks in Large, Dynamic Robot Teams (short paper) Proceedings of the 2011 Conference on Autonomous Agents and Multi-Agent Systems, May 2-6, Taipei, Taiwan, 2011.
114. Saqer, H., de Visser, E., Emfield, A., Shaw, T., & Parasuraman, R. (2011). Adaptive automation to improve human performance in supervision of multiple uninhabited aerial vehicles: Individual markers of performance. In Proceedings of the Annual Conference of the Human Factors and Ergonomics Society, Santa Monica, CA.
115. Ponda, S., Ahmed, N., Luders, B., Sample, E., Levine, D., Hoossainy, T., Shah, D., Campbell, M., and How, J. P., "Decentralized Information-Rich Planning and Hybrid Sensor Fusion for Uncertainty Reduction in Human-Robot Missions," AIAA Guidance, Navigation and Control Conference, (GNC), Portland, OR, August 2011 (**Best paper award**).

116. N. Ahmed, M. Campbell "Variational Learning of Autoregressive Mixture of Experts Models for Fully Bayesian Non-Markov Hybrid System Identification," 2011 American Control Conference.
117. N. Ahmed, E. Sample, K. Ho, T. Hoossainy, M. Campbell "Categorical Soft Data Fusion for Gaussian Mixtures via Variational Bayesian Importance Sampling, with Applications to Cooperative Search," 2011 American Control Conference.
118. D. Shah, M. Campbell, "A Robust Qualitative Planner for Mobile Robot Navigation Using Human-Provided Maps," 2011 International Conference on Robotics and Automation.
119. Tinglong Dai, T., Sycara, K. & Lewis, M. (2011) A game theoretic queuing approach to self-reflection in decentralized human-robot interaction systems, 2011 IEEE International Conference on Robotics and Automation
120. Wang, H., Kolling, A., Abedin, S., Lee, P., Chien, S., Lewis, M., Brooks, N., Owens, S., Scerri, P. & Sycara, K. (2011) Scalable target detection for large robot teams, Proceedings of the 6th ACM/IEEE International Conference on Human-Robot Interaction.
121. A.S. Clare and M.L. Cummings, Task-Based Interfaces for Decentralized Multiple Unmanned Vehicle Control, Proceedings of AUVSI 2011: Unmanned Systems North America, Washington D.C., August 2011.
122. Lee, P., Kolling, A., & Lewis, M. (2011). Workload Modeling using Time Windows and Utilization in an Air Traffic Control Task, Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11).
123. Glinton, R., Sycara, K., Scerri, P. "Exploiting Scale Invariant Dynamics for Efficient Information propagation in Large Teams", Proceedings of the 2010 Conference on Autonomous Agents and Multi-Agent Systems, Toronto, CA, May, 2010 (**Second Place for Best Paper Award**).
124. Lisy V., Zivan R., Sycara, K. Pechoucek M. "Deception in Networks of Mobile Sensing Agents", Proceedings of the 2010 Conference on Autonomous Agents and Multi-Agent Systems, Toronto, CA, May, 2010.
125. Chakraborty, N., Sycara, K. "Reconfiguration Algorithms for Mobile Robotic Networks", ICRA, Alaska, 2010.
126. Kolling, A., Kleiner, A., Lewis, M., Sycara K. "Pursuit-Evasion in 2.5d Based on Team Visibility", Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10), October 18-22, Taipei, Taiwan, 2010.
127. Velagapudi, P., Sycara, K., Scerri, P. "Decentralized Prioritized Planning in Large Multirobot Teams", Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10), October 18-22, Taipei, Taiwan, 2010.
128. Xu, Y., Dai, T., Sycara, K., Lewis, M. "Service Level Differentiation in Multi-Robot Teams", IROS, October 18-22, Taipei, Taiwan, 2010.
129. Scerri, P., Velagapudi, P., Sycara, K., Wang, H., Chien, S. & Lewis, M. Towards an understanding of the impact of autonomous path planning on victim search in USAR, Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10), October 18-22, 2010, 383-388.
130. Lee, P., Wang, H., Chien, S., Lewis, M., Scerri, P., Velagapudi, P., Sycara, K. & Kane, B. Teams for Teams: Performance in Multi-Human/Multi-Robot Teams.

- Proceedings of the 54th Annual Meeting Human Factors and Ergonomics Society (HFES'10), September 27- October 1, 2010.
131. Scerri, P., Owens, S., Sycara, K., Lewis, M. User evaluation of a GUI for controlling autonomous persistent surveillance team, In SPIE 2010.
  132. Wang, H., Lewis, M., Chien, S., Scerri, P., Velagapudi, P., Sycara, K. & Kane, B. (2010). Teams organization and performance in multi-human/multi-robot teams, 2010 IEEE International Conference on Systems, Man, and Cybernetics, (SMC'10), October 10-13.
  133. Glinton, R., Scerri, P., Sycara, K. "An Explanation for the Efficiency of Scale Invariant Dynamics of Information Fusion in Large Teams", International conference on Information Fusion (Fusion2010), July 26-29, Edinburgh, UK, 2010.
  134. L. Luo, N. Chakraborty, and K. Sycara, "Prisoner's Dilemma in Graphs with Heterogeneous Agents," IEEE International Conference on Social Computing, 2010, Minneapolis, MN.
  135. L. Luo, N. Chakraborty, and K. Sycara, "Modeling effect of leaders in ethno-religious conflicts," 2010 International Conference on Social Computing, Behavioral Modeling, and Prediction (SBP10), Washington D.C., USA, March 2010
  136. Zivan, R., Dudik, M., Okamoto S., Sycara K. "Reducing Untruthful Manipulation in Envy-Free Pareto Optimal Resource Allocation", International Conference on Intelligent Agent Technology, Aug 30-Sep 3, Toronto, CA, 2010.
  137. Paruchuri, P., Varakantham, P., Sycara, K., Scerri, P. "Effect of Human Biases on Human Agent Teams", International Conference on Intelligent Agent Technology, Aug 30-Sep 3, Toronto, CA, 2010.
  138. Xu, Y., Dai, T., Sycara, K. & Lewis, M. (2010). Service level differentiation in multi-robots control, Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10), October 18-22, Taipei, Taiwan, 2224-2230.
  139. Chien, S., Wang, H., & Lewis, M. (2010). Human vs. algorithmic path planning for search and rescue by robot teams, Proceedings of the 54th Annual Meeting of the Human Factors and Ergonomics Society (HFES'10), 379-383.
  140. Scerri, P., Owens, S., Sycara, K. & Lewis, M. (2010). User evaluation of a GUI for controlling an autonomous persistent surveillance team, In SPIE'10.
  141. Kolling, A., Kleiner, A., Lewis, M., & Sycara, M. (2010). Pursuit-evasion in 2.5d based on team-visibility, Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10), 4610 – 4616.
  142. N. Ahmed, M. Campbell, "Variational Bayesian Data Fusion of Multi-class Discrete Observations with Applications to Cooperative Human-Robot Estimation," 2010 International Conference on Robotics and Automation.
  143. D. Lee, et al, "Distributed, Collaborative Human-Robotic Networks for Outdoor Experiments in Search, Identify and Track," 2010 SPIE Europe Conference on Unmanned/Unattended Sensors and Sensor Networks.
  144. D. C. Shah, J. R. Schneider, M. E. Campbell, "A Robust Sketch Interface for Natural Robot Control," submitted to the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems.

145. S. Ponda, J. Redding, H.L. Choi, J.P. How, M. Vavrina, J. Vian, "Decentralized Planning for Complex Missions with Dynamic Communication Constraints", American Control Conference, 2010, Baltimore, MD
146. S. Ponda, H.L. Choi, J.P. How, "Predictive Planning for Heterogeneous Human-Robot Teams", AIAA Infotech@Aerospace, 2010, Atlanta, GA
147. L. Johnson, S. Ponda, H.L. Choi, J.P. How, "Improving the Efficiency of a Decentralized Tasking Algorithm for UAV Teams with Asynchronous Communications", AIAA Guidance, Navigation, and Control Conference, 2010, Toronto, Canada
148. H.L. Choi, A. Whitten, J.P. How, "Decentralized Task Allocation for Heterogeneous Teams with Cooperation Constraints", American Control Conference, 2010, Baltimore, MD
149. Lebiere, C., Stocco, A., Reitter, D., & Juvina, I. (2010). Scaling Up High-Fidelity Cognitive Modeling to Real-World Applications. In *Proceedings of NATO Workshop on Human Modeling for Military Application*. Amsterdam, NL: October 18-20, 2010.
150. Reitter, D., & Lebiere, C. (2010). Accountable Modeling in ACT-UP, a Scalable, Rapid-Prototyping ACT-R Implementation. In *Proceedings of the 2010 International Conference on Cognitive Modeling*. Philadelphia, PA.
151. Reitter, D., Juvina, I., Stocco, A., & Lebiere, C. (2010). Resistance is Futile: Winning Lemonade Market Share through Metacognitive Reasoning in a Three-Agent Cooperative Game. In *Proceedings of the Behavior Representation In Modeling and Simulations (BRIMS 2010) Conference*. Charleston, SC.
152. Reitter, D., & Lebiere, C. (2010). Did social networks shape language evolution? a multi-agent cognitive simulation. In *Proc. Cognitive Modeling and Computational Linguistics Workshop (CMCL 2010, at Association for Computational Linguistics: ACL 2010)*, Uppsala, Sweden.
153. Reitter, D., & Lebiere, C. (2010). On the influence of network structure on language evolution. In R. Sun, editor, *Proc. Workshop on Cognitive Social Sciences: Grounding the Social Sciences in the Cognitive Sciences (at Cognitive Science: CogSci 2010)*, Portland, Oregon.
154. McDonald, A.D., (2010) A Discrete Event Simulation Model For Unstructured Supervisory Control Of Unmanned Vehicles, B. S. Thesis, MIT Mechanical Engineering, Cambridge, MA.
155. Parasuraman, R., Mohammed-Ameen, A., de Visser, E., & Shaw, T. H (2010). Verification behavior with imperfect automated decision aids in networked systems. In *Proceedings of the Human Factors and Ergonomics Society Conference*.
156. de Visser, E., Shaw, T., Mohamed-Ameen, A., & Parasuraman, R. (2010,). Modeling human-automation team performance in networked systems: Individual differences in working memory count. *Proceedings of the Human Factors and Ergonomics Society, Santa Monica, CA: Human Factors and Ergonomics Society*.
157. Shaw, T., Guagliardo, L., de Visser, E., & Parasuraman, R. (2010). Using Transcranial Doppler Sonography to measure cognitive load in a command and control task. *Proceedings of the Human Factors and Ergonomics Society*, 249-253. Santa Monica, CA: Human Factors and Ergonomics Society.

158. Squire, P. N., & Parasuraman, R. (2010). Effects of automation and task load on task switching during human supervision of multiple semiautonomous robots in a dynamic environment. *Ergonomics*, 53(8), 951-961.
159. Velagapudi, P., Prokopyev, O., Sycara, K., Scerri, P. Analyzing the performance of randomized information sharing, AAMAS 2009, Budapest, Hungary, May 15-19, 2009.
160. Owens, S., Sycara, K., Scerri, P. "Using Immersive 3D Terrain Models for Fusion of UAV Surveillance Imagery", In Proceedings of AIAA, 2009.
161. Glinton, R., Scerri, P., Sycara, K. Towards the Understanding of Information Dynamics in Large Scale Networked Systems, Twelfth International conference on information Fusion 2009, Seattle WA., July 6-9, 2009.
162. Glinton, Scerri, Paruchuri, Sycara "Analysis and Design of Information Dynamics in Large Scale Networked Systems" International Conference on the Dynamics of Information System, Gainesville, Fla, Jan 28-30, 2009.
163. Paruchuri, Glinton, Scerri, Sycara "The role of humans in information propagation in Large Scale Heterogeneous Teams" International Conference on the Dynamics of Information Systems, Gainesville, Fla, Jan 28-30, 2009
164. Velagapudi, P., Owens, S., Scerri, P., Sycara, K., Lewis, M., "Environmental Factors Affecting Situation Awareness in Unmanned Aerial Vehicles", In Proceedings of AIAA, 2009.
165. Lingzhi Luo, Nilanjan Chakraborty, and Katia Sycara. Prisoner's dilemma on graphs with heterogeneous agents. Proceedings of the 11th annual conference companion on Genetic and evolutionary computation conference (GECCO09). Montreal, Canada, Jul. 2009.
166. Lewis, M., Wang, J., Velagapudi, P., Scerri, P., Sycara, K. "Using Humans as Sensors in Robotic Search", Proceedings of the 12th International Conference on Information fusion, Seattle WA., July 6-9, 2009.
167. Velagapudi, P., Wang, H., Scerri, P., Lewis, M., Sycara, K. "Scaling Effects for Streaming Videos vs. Static panorama in Multi-Robot Search", International conference on Intelligent Robots and Systems (IROS'09), IEEE, 2009.
168. Wang, H., Lewis, M., Velagapudi, P., Scerri, P., Sycara, K. "How Search and its Subtasks Scale in N Robots", Proceedings of the ACM/IEEE International conference on Human Robot Interaction (HRI'09), IEEE, pp. 141-147, 2009.
169. Wang, H. Chien, S. , Lewis, M., Valagapudi, P., Scerri, P., Sycara, K. "Human Teams for Large Scale Multirobot Control", proceedings of the 2009 IEEE International Conference on Systems, Man and Cybernetics, pp. 1306-1311, 2009.
170. Wang, H., Lewis, M., Velagapudi, P., Scerri, P., and Sycara, K. (2009). Scaling effects for synchronous vs. asynchronous video in multi-robot search, Proceedings of the 53rd Annual Meeting of the Human Factors and Ergonomics Society, October.
171. Zivan, R., Glinton, R, Sycara, K., Distributed Constraint Optimization for large teams of mobile sensing agents, In proceedings of the International conference on Intelligent Agent Technology, Milan, Italy, September 15-18, 2009.
172. Lewis, M. & Wang, J. (2009). Measuring coordination demand in multirobot teams, Proceedings of the Human Factors and Ergonomics Society 53rd Annual Meeting (HFES'09), 779-783.

173. Ma, Z., Lee, P., Xu, Y., Lewis, M. & Scerri, P. (2009). An initial evaluation of approaches to building entry for large robot teams, 2009 IEEE International Workshop on Safety, Security and Rescue Robotics (SSRR'09), IEEE.
174. Reitter, D., Lebiere, C., Lewis, M., Wang, H. & Ma, Z. (2009). A cognitive model of perceptual path planning in a multi-robot control system, Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics, 2976-2982.
175. Reitter, D., & Lebiere, C. (2009). Towards Explaining the Evolution of Domain Languages with Cognitive Simulation. In Proceedings of the 9th International Conference on Cognitive Modeling. Manchester, England.
176. Mekdeci, B., and Cummings, M. L. (2009). Modeling Multiple Human Operators in the Supervisory Control of Heterogeneous Unmanned Vehicles, 9th Conference on Performance Metrics for Intelligent Systems (PerMIS'09), Gaithersburg, MD, September 21-23, 2009,
177. Baishya, D. & Lewis, M. (2009). Algorithm steering for mixed-initiative robot teams, 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'09) Workshop on Mixed-Initiative MAS, May 10-15.
178. Nehme, C. "Modeling human supervisory control in heterogeneous unmanned vehicle systems," Ph.D. dissertation, Massachusetts Institute of Technology, Dept. of Aeronautics and Astronautics, 2009.
179. Parasuraman, R. (2009). Neuroergonomics applied to adaptive automation for supervision of multiple unmanned vehicles. In Proceedings of the Army Science Conference, Orlando, FL.
180. Reitter D. and Lebiere C. A subsymbolic and visual model of spatial path planning. In: Proc. Behavior Representation in Modeling and Simulation (BRIMS), 2009. Best paper award BRIMS 2009.
181. Reitter, D. Lebiere, C. Lewis, M. Wang, H. and Ma Z.. A cognitive model of visual path planning in a multi-robot control system. In: Proceedings Systems Man Cybernetics 2009 (IEEE-SMC), San Antonio, TX, October 2009.
182. Undurti A and How, J. "Task planning for teams with disobedient agents," IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009.
183. Velagapudi, P., Scerri, P., Sycara, K., Wang, H., Lewis, M., and Wang, J. (2008a). Scaling effects in multi-robot control, 2008 International Conference on Intelligent Robots and Systems (IROS'08).
184. Velagapudi, P., Wang, H., Scerri, P., Lewis, M., and Sycara, K. (2008b). Potential scaling effects for asynchronous video in multirobot search, Performance Metrics for Intelligent Systems (PerMIS'08).
185. Grier, R. A., Parasuraman, R., Entin, E. E., Bailey, N., & Stelzer, E. (2008). A test of intra- versus inter-modality interference as a function of time pressure in a warfighting simulation. In Proceedings of the Human Factors and Ergonomics Society. Santa Monica, CA, pp, 1229-1232.
186. Nehme, C. Kilgore, R and Cummings, M. "Predicting the impact of heterogeneity on unmanned vehicle team performance," in Proceedings of HFES 2008: 52nd Annual Meeting of the Human Factors and Ergonomic Society, New York, NY, USA, September 2008.

187. Giampapa, J. Sycara, K., and Sukthankar, G. "Toward Identifying Process Models in Ad Hoc and Distributed Teams," Proceedings of the First International Working Conference on Human Factors and Computational Models in Negotiation (HuCom 2008), Delft, the Netherlands December 8-9, 2008, pp. 55-62.

**Interactions (eg presentations at Dod meetings, keynote talks, panels, seminars etc)**

Katia Sycara (Keynote Talk) “Scalable Strategies of Human Control for Multi Robot Systems”, The 16<sup>th</sup> World Scientific and Engineering Academy and Society on Computers, Kos, Greece, July 14-14, 2012.

Katia Sycara (Keynote Talk) “Network Dynamics of Information Propagation” 5<sup>th</sup> International Symposium on Intelligent Distributed Computing, Delft, the Netherlands, October 5-7, 2011.

Missy Cummings gave an invited presentation at the The U.S. Air Force Tactical Mission Battlespace Integration workshop, Alexandria, VA, May 2011

Missy Cummings briefed The Office of the Secretary of Defense, Washington DC in December 2010.

Missy Cummings was the keynote speaker for the 2011 Naturalistic Decision Making conference JUN 2011

Missy Cummings gave presentations for Codes 34 & 35 (JUN 2011 and APR 2011 respectively).

Raja Parasuraman made presentations at the Human Effectiveness Directorate, Air Force Research Laboratory, Wright Patterson Air Force Base in 2011 and 2012.

Mark Campbell gave an Invited Talk, National Academy of Science Chinese-American Kavli Frontiers of Science Symposium, Shenzhen, China, Nov 2011.

Mark Campbell Plenary Talk, NSF-ARO Frontiers of Real-World Multi-Robot Systems: Challenges and Opportunities, 10 Oct 2011.

Mark Campbell gave a Plenary Talk, NSF-ARO Frontiers of Real-World Multi-Robot Systems: Challenges and Opportunities

David Reitter attended BRIMS-2012 (sponsored by ARL) in Florida in March 2012, including a workgroup with Capt Dylan Schmorow, Deputy Director, Human Performance, Training and BioSystems at the Office of the Secretary of Defense.

Christian Lebiere gave an invited talk at the University of Indiana School of Public and Environmental Affairs in May 2012 on the topic of “Cognitive Architectures for Social Decision Making” emphasizing our multi-agent modeling approach.

Christian Lebiere presented a presentation titled “Introduction to ACT-R Cognitive Architecture for Robotics” at an ONR-sponsored workshop on Architectures for Autonomy in July 2012 in Arlington, VA.

### **Consultative and advisory functions at other laboratories and agencies, especially Air Force and Other DoD laboratories.**

Michael Lewis: Invited Participant in Human-Centered Autonomy Workshop at WPAFB. He was member of the group reporting on research needs for Human Robotics Interface, September 13-14, 2011.

Michael Lewis was an invited Participant in HFM-217 NATO Workshop on Supervisory Control of Multiple Uninhabited Systems - Methodologies and Human-Robot Interface Technologies, Prague, Czech Republic, May 8-10 2012.

Raja Parasuraman was invited to Chair the Panel on Sensing, Air Force sponsored Workshop on Human Performance Augmentation, Arizona State University, Tempe, AZ, 2012.

Parasuraman: Human Effectiveness Directorate, Air Force Research Laboratory, Wright Patterson Air Force Base, Dayton, OH. Collaborative research on multi-UAV supervisory control, adaptive automation, and neuroergonomics.

Parasuraman: United States Military Academy, West Point, NY. Collaborative research on field study of individual differences in human performance, with West Point Marine cadets.

Parasuraman was consulted by Air Force Research Laboratory, Human Effectiveness Directorate on issues related to multi-UAV control, adaptive automation, and neuroergonomics.

Lebiere gave input at a workshop on a tri-service modeling competition organized by Research Psychologist Kevin Gluck from AFRL Dayton, held at Aberdeen Proving Ground, and submitted a whitepaper about topics and challenges for future competitions.

### **Interactions/Transitions:**

The CMU multi-robot planning POMDP algorithm is being transitioned to an ONR HSBC program in a project called “Enhanced COA Analysis by Integration of Decision and Social Influence Modeling with MultiAgent System Technology (CADSIM)” in cooperation with a small government contractor. The algorithms will form the basis of a war-gaming and contingency planning system.



The CMU multi-robot path planning POMDP algorithm is also being brought into an AFOSR SBIR program as a planning service for operators controlling multiple UAVs.

The CMU information sharing algorithms for large scale networks are being adapted, in part, into an Army SBIR for extracting information from very large databases and the web and automatically constructing networks of interactions between people, places and organizations.

The CMU ACT-UP cognitive modeling toolbox is available online for download at:

<http://act-up.psy.cmu.edu/ACT-UP.html>  
<https://cc.ist.psu.edu/act-up/>

To facilitate its adoption by the cognitive modeling community, supporting materials such as tutorial examples, documentation, and mailing list are also provided.

**Patent disclosures:**

--None

## Appendix A: ACT-UP Documentation

### Overview

ACT-UP is a cognitive modeling library that allows modelers to specify their model's functionality in Common Lisp. Whenever a cognitive explanation in a particular part of the model is sought, the modeler uses the library to provide characteristics of

- explicit, declarative learning and cue- and similarity-based retrieval, and
- procedural skill acquisition [not yet available]

following the ACT-R 6 theory.

As in the ACT-R 6 implementation, modelers are free to adhere more or less to the theoretical limitations. However, ACT-UP's design encourages modelers to underspecify portions of the model's functionality that do not contribute to the model's explanations and predictions of human performance.

### How do I...

*... load the library?*

Just load the file `load-act-up.lisp`: The easiest way is to store the ACT-UP directory somewhere on your hard drive and then hard-code the path:

```
(load "/Users/me/modeling/ACT-UP/load-act-up.lisp")
```

**Windows users, beware: backslashes need to be doubled in Lisp strings; forward slashes should work fine.**

A more sophisticated solution uses a path relative to the model file. Assuming our model file is `ACT-UP/tutorials/model.lisp`, do this:

```
(load (concatenate 'string (directory-namestring *load-true-name*) "../load-act-up.lisp"))
```

Here, we adjust its path so that it is relative to the current file (rather than the directory that happens to be current when Lisp is started or when the model file is loaded).

*... define a chunk type?*

Unlike ACT-R, ACT-UP is not normally strongly typed. All slot names are declared initially, but ACT-UP does not distinguish chunk types within a type hierarchy. Chunk types are

lisp structure types that inherit from the type [actup-chunk](#). This type exists once the 'define-slots' macro is called:

```
(define-slots name dampen success)
```

If you do want to define a type hierarchy, ACT-UP provides the necessary macros. For example, the following structure defines a chunk type of name [strategy](#) with four slots. One of these slots is assigned a default value ([strategy](#)).

```
(define-chunk-type strategy
  (type 'a-strategy)
  name
  dampen
  success)
```

Note that the [type](#) member is not required by ACT-UP.

To define an inherited type, use this construction:

```
(define-chunk-type (lazy-strategy :include strategy)
```

*...define a new model?*

The model is defined automatically when [act-up.lisp](#) is loaded. To reset the model, use the [reset-model](#) function. To create a new model (multiple models may be used in parallel), use [make-model](#). Use the function [set-current-actUP-model](#) to define the current model.

The ACT-UP meta-process keeps track of model time that is common to all models. You may define several meta-processes and use/reuse them as you like with the function [make-meta-process](#). You can bind `*current-actUP-meta-process*` to a meta-process to switch. Use [reset-mp](#) to discard and reset the current meta-process.

*... define a procedural rule ("production")?*

ACT-UP does not use IF-THEN production rules as known from ACT-R. Instead, it allows you define Lisp functions that we call *procedures*; they represent multiple, theoretical ACT-R productions. An important property of ACT-UP models is that the procedures are not always tested in parallel; flow control is achieved through standard Lisp programming. Define procedures using [defproc](#), similar to the way you would define a Lisp function with the 'defun' macro:

```
(defproc subtract-digit (minuend subtrahend)
  "Perform subtraction of a single digit"
  (- minuend subtrahend))
```

*...define a chunk?*

Chunks are Lisp structures that are of type `'chunk'`, or of a type defined with `'define-chunk-type'`. They can be created with the `'make-chunk'` function, or with the creator functions of the more specific type.

When a chunk is created, a unique name should be assigned. Otherwise, this name is assigned automatically when the chunk is added to the DM.

```
(make-chunk :name 'andrew :age 42 :spouse 'louise)
(make-chunk :name 'louise :age 35 :spouse 'andrew)
```

When assigning values to the attributes defining a chunk, symbols are interpreted as names of other chunks in DM. This is often more comfortable than assigning the values directly.

Note, however, that certain actions - such as defining Sji weights between chunks - will cause ACT-UP to implicitly define an empty chunk of a given name in the DM, if that chunk is not found in DM.

*...commit a chunk to memory or reinforce it?*

To specify the "presentation" of a specific chunk, use the function [learn](#). The chunk reference may be supplied in a normal variable (equivalent to ACT-R's buffer), or the chunk may be produced right there and then using the [make-chunk](#) function, as in the following example:

```
(learn-chunk (make-chunk :name 'guess :success 0.2))
```

This will create a new strategy chunk, setting two of its parameters, and commit it to memory. To reinforce the existing chunk, use the chunk's name:

```
(learn-chunk 'guess)
```

Note that making a new chunk and calling `'learn-chunk'` will always create a separate chunk. It will not merge the new chunk with any existing chunk (this would not scale very well, computationally). You must use the unique chunk name, or retrieve the chunk before committing it, or use the `'make-chunk*'` syntax to extract a chunk from declarative memory for learning. For example:

```
(learn-chunk (make-chunk :success 0.2)) [1]
(learn-chunk (make-chunk* :success 0.2)) [2]
```

Case 1 would make a new chunk with the given success value, give it a unique name, and add it to declarative memory. Case 2, on the other hand, would find the chunk that is already in declarative memory, and boost its presentation count via base-level learning.

... retrieve an item from declarative memory?

Simply use the high-level functions [retrieve-chunk](#), or [blend-retrieve-chunk](#) (for blending). The following example retrieves the most active chunk that has the name [guess](#). The chunk contained in the variable *valve-open-chunk* spreads activation. No partial matching is used:

```
(retrieve-chunk '(:name guess))  
:cues (list valve-open-chunk)
```

Several low-level functions are provided as well. [filter-chunks](#) produces a list of all chunks that match a given set of criteria. In the example below, we are looking for a chunk with the *name* attribute [guess](#).

The [best-chunk](#) function does the actual (time-consuming and noisy) retrieval: it selects the best chunk out of the (filtered) list of chunks, given additional retrieval cues that spread activation and, if so desired, a set of filter specifications for partial matching. In this example, we use an existing chunk stored in the *valve-open-chunk* variable as a single retrieval cue, and no partial matching:

```
(best-chunk (filter-chunks  
            (model-chunks (current-actUP-model))  
            '(:name guess))  
            (list valve-open-chunk)  
            nil)
```

... debug an ACT-UP model? ("production")?

We're providing a separate tutorial on [debugging ACT-UP models](#).

... retrieve a blended chunk?

Use the high-level function [blend-retrieve-chunk](#).

When combining low-level functions, use the function [blend](#) instead of [retrieve-chunk](#). In addition to the cues and partial-matching specification known from [retrieve-chunk](#), it also expects a chunk type (such as [strategy](#)), which determines the kind of chunk created as a result of blending.

... define chunk similarities?

Use the [add-sji-fct](#) and `reset-sji-fct'` functions.

... select a procedure (in lieu of a production rule) using subsymbolic utility learning?

### Quick Answer

Define competing procedures as above and give each a `:group` attribute in order to group them into a competition set:

```
(defproc force-over ()
  :group choose-strategy
  ...)
(defproc force-under ()
  :group choose-strategy
  ...)
```

Then, invoke one of the procedures (as chosen by utility) as such:

```
(choose-strategy)
```

Arguments may be used as well (but ensure that all procedures accept the same arguments).

Utilities are learned using the function [assign-reward](#):

```
(assign-reward 1.5)
```

This example distributes a reward of across the recently invoked procedures. Procedures do not have to have a `:group` attribute and they do not have to have been invoked via the group name in order to receive a reward; however, they have to have been defined using the [defproc](#) macro (rather than just being Lisp functions).

Configure utility learning via the parameters [\\*au-rpps\\*](#), [\\*au-rfr\\*](#), [\\*alpha\\*](#), and [\\*iu\\*](#).

### Worked Example

Note that ACT-UP supports utility learning and even procedure compilation. Utility learning means that multiple procedures may compete for execution, and that the actually executed procedures are assigned rewards if they lead to some form of success. To define competing procedures, they must be grouped together in a *Group*. A group is a set of procedures, such as the following:

```
(defproc subtract-digit-by-addition (minuend subtrahend)
  :group subtract
  "Perform subtraction of a single digit via addition."
  (let ((chunk (retrieve-chunk `(:chunk-type addition-fact
                                :result ,minuend
                                :add1 ,subtrahend))))
    (when chunk
      (learn-chunk chunk)
      (addition-fact-add2 chunk))))
(defproc subtract-digit-by-subtraction (minuend subtrahend)
```

```

:group subtract
"Perform subtraction of a single digit via subtraction knowledge."
(let ((chunk (retrieve-chunk `(:chunk-type addition-fact
                               :min ,minuend
                               :sub subtrahend))))
  (print "addition by subtraction.")
  (when chunk
    (learn-chunk chunk)
    (subtraction-fact-result chunk))))
(defproc subtract-digit-by-addition-faulty (minuend subtrahend)
:group subtract
"Perform subtraction of a single digit via addition. Faulty
strategy."
(let ((chunk (retrieve-chunk `(:chunk-type addition-fact
                               :add2 ,minuend
                               :result ,subtrahend))))
  (when chunk
    (learn-chunk chunk)
    (addition-fact-add2 chunk))))
(defproc subtract-digit-by-decrement (minuend subtrahend)
:group subtract
"Perform subtraction of a single digit via subtraction knowledge."
...)

```

Note that each procedure in the group takes the same, two arguments (minuend, subtrahend). In order to execute a subtraction, we simply call a function that is named after the group:

```
(subtract 7 3)
```

ACT-UP will automatically choose one of procedures in the *subtract* group. In order to gauge the utility of each group, we must propagate rewards to the procedures. This can be done with the 'assign-reward' function:

```

(defproc subtraction-model (a b)
  (let ((result (subtract a b)))
    ;; obtain feedback from experimental environment:
    (if (get-feedback a b result)
        (assign-reward 2.0))))

;; environment:
(defun get-feedback (a b result)
  "Environment function (experimental setup) - not part of the model.
Return T if problem solved correctly."
  (if result ;; note result may be nil
      (= result (- a b))))

```

After a short period of time, this model should learn to choose an effective, reliable strategy to carry out a subtraction.

Rewards are assigned to ACT-UP procedures just like they would be assigned to production procedures in ACT-R. The most recently invoked procedure received the largest portion of the reward; Difference-Learning governs how much of a procedure benefits from its reward portion.

*...model effects via production compilation?*

ACT-UP may not have productions, but it does have *procedures*. These procedures can be compiled. To do so, we need to keep in mind that procedure compilation will side-step any intermediate action that a model might undertake in order to execute a procedure. This includes retrievals from declarative memory, but also any other side-effects such as sensory-motor interaction, or even Lisp code.

ACT-UP's procedure compilation can be enabled by setting the `*procedure-compilation*` parameter to `t`.

Every time a procedure ("source procedure") is invoked, procedure compilation will create a compiled procedure (specific to the arguments given to the procedure at invocation); its initial utility will be `*iu*`. If the source procedure is compiled a second time, the utility of the compiled procedure will be boosted by the utility of the source procedure according to reward assignment mechanism (difference learning equation, as above). The compiled procedure will, eventually, have a higher utility than the source procedure, and it will be executed instead. In the subtraction example from above, the following gives a sample of the acquired compiled procedures:

```
(subtract-digit-by-subtraction 8 3) --> 5 (subtract-digit-by-addition 6 2) --> 4 (subtract-digit-by-addition-faulty 7 3) --> nil (subtraction-model 3 1) --> 2
```

Note that in order to execute the best procedure among one or more source procedures, and all their compiled equivalents, the modeler must define a group for the procedures and invoke them via the group name. Reward assignment and procedure compilation will take place no matter how the procedure was invoked. So, in the above example, the `'subtraction-model'` will never be run in its compiled form, and rewards will always be propagated, because it does not belong to a group and cannot be called that way.

Again, note that in its compiled form, the procedure merely returns its result. No side-effects are observed. For instance, the `'subtract-digit-by-subtraction'` procedure will print the debug message "Addition by subtraction!" every time it is run - as long as it isn't compiled. Once compiled, it will always just return the result.

*... run model code in parallel?*

ACT-UP is designed assuming that most modeled processes can be formulated as a sequence of cognitive actions. However, in some situations, parallelism may be necessary.

To asynchronously request the execution of some code (that is, without waiting for the results), use the `'request-'` syntax, e.g., `request-retrieve-chunk`. The *request-* functions are defined for each module-specific ACT-UP function that can take some time, e.g., `best-chunk`, `filter-chunks`, `retrieve-chunk` (for the declarative memory module), and



all functions defined with ``defproc'` (for the procedural module). The functions all return an execution handle.

This function kicks off task execution in parallel; it returns without delay (in ACT-UP time). Once the result of the operation is needed, it may be retrieved using the ``receive'` function and the previously obtained *handle*.

```
(let ((handle (request-retrieve-chunk '(:chunk-type ...))))  
;; do something else  
...  
(receive handle))
```

Different threads of execution may share resources. We follow Anderson et al. (2004) in that each module can only handle one request at a time. We follow some of Salvucci&Taatgen's (2008) *threaded cognition* approach: threads acquire resources in a "greedy" and "polite" manner. When a ``request-'` function is called, it will wait until the module is available, but then reserve the module regardless of other goals that may exist. The module functions (such as ``retrieve-chunk'`) will also wait for the module to be free. Similarly, ``receive'` will wait. To check if the result is available, use the ``response-available-p'` function.

### Example

The following example shows how a retrieval request is initiated and finished. Upon initiating the request, ACT-UP does not "wait" for the retrieval to finish.

```
(print (actup-time))  
(let ((retrieval-process (request-retrieve-chunk '(:chunk-type person))))  
  (print (actup-time)) ;; no time has elapsed  
  (print (response-available-p retrieval-process)) ;; module is busy  
  (pass-time 0.05) ;; let's spend some time  
  (print (response-available-p retrieval-process)) ;; module is still busy  
  ;; (wait-for-response retrieval-process) ;; wait for result - not needed  
  ;; (print (response-available-p retrieval-process))  
  (print (actup-time)) ;; this takes some time!  
  (print (receive retrieval-process)) ;; waits and receives
```

### More related functions

ACT-UP provides a ``reset-module'` function to explicitly cancel a module's operation. To wait for a module to finish processing when the handle is not known, use ``wait-for-module'`.

## Appendix B: ACT-UP Package API

*\*act-up-version\* variable*

Version of a loaded ACT-UP. ACT-UP has been correctly initialized if this is defined and non-nil.

Initial value: "27bc8ed"

*\*all\* variable*

Constant for [\\*debug\\*](#): Show all messages (maximum detail).

Initial value: 1000

*\*alpha\* variable*

Utility learning rate. See also the function [assign-reward](#). See also: ACT-R parameter :alpha

Initial value: 0.2

*\*ans\* variable*

Transient noise parameter for declarative memory. See also: ACT-R parameter :ans

Initial value: 0.2

*\*associative-learning\* variable*

The trigger for associative learning,  $a$  in ROM Equation 4.5.

Can be any non-negative value.

Initial value: NIL

*\*au-rfr\* variable*

base reward proportion for each procedure e.g., the each procedure before the reward trigger gets 10% of the reward. Set to nil (default) to use the ACT-R discounting by time in seconds. See also the parameter [\\*au-rpps\\*](#) and the function [assign-reward](#).

Initial value: NIL

*\*au-rpps\* variable*

Reward proportion per second elapsed. e.g., after 10 seconds we want to assign 50% of the remaining reward:  $0.5/10 = 0.05$  time is in between procedures. Set to nil (default) to use the ACT-R discounting by time in seconds. See also the parameter [\\*au-rfr\\*](#) and the function [assign-reward](#).

Initial value: NIL

*\*blc\* variable*

Base-level constant parameter for declarative memory. See also: ACT-R parameter :blc  
Initial value: 0.0

*\*bll\* variable*

Base-level learning decay parameter for declarative memory. See also: ACT-R parameter :bll  
Initial value: 0.5

*\*critical\* variable*

Constant for [\\*debug\\*](#): Show only critical messages.

Initial value: 0

*\*current-actup-meta-process\* variable*

The current ACT-UP meta-process. The meta process keeps track of simulation time. May be read and manipulated by setting it to a different instance of type [meta-process](#).

Initial value: #S(META-PROCESS :ACTUP-TIME 0.0D0 :NAME NIL)

*\*dat\* variable*

Default time that it takes to execut an ACT-UP procedure in seconds. See also: ACT-R parameter :dat [which pertains to ACT-R productions]

Initial value: 0.05D0

*\*debug\* variable*

Level of debug output currently in effect. The following constants may be used: *\*critical\** *\*warning\** *\*informational\** *\*all\** The parameter [\\*debug-to-log\\*](#) is helpful in logging debug messages to a file.

Initial value: 10

*\*debug-to-log\* variable*

Enable off-screen logging of debug output. If t, ACT-UP logs all debug messages not to standard output, but to a buffer that can be read with [debug-log](#) and cleared with [debug-clear](#). If a stream, ACT-UP logs to the stream.

Initial value: NIL

*\*declarative-finst-span\* variable*

Declarative Finst time span The maximum time period during whichg a finst marks a chunk as recently retrieved. Chunks retrieved longer ago are not considered 'recently retrieved'. Time in seconds, defaults to 3.0. See ACT-R parameter :declarative-finst-span

Initial value: 3.0

*\*declarative-num-finsts\* variable*

Number of Declarative Finsts The maximum number of chunks considered recently retrieved. Defaults to 4. See ACT-R parameter :declarative-num-finsts

Initial value: 4

#### *\*detailed\* variable*

Constant for [\\*debug\\*](#): Show detailed log output .

Initial value: 300

#### *\*egs\* variable*

Transient noise parameter for ACT-UP procedures. This is the expected gain s parameter. It specifies the s parameter for the noise added to the utility values. It defaults to 0 which means there is no noise in utilities. See also: ACT-R parameter :egs

Initial value: NIL

#### *\*informational\* variable*

Constant for [\\*debug\\*](#): Show informational and more important messages.

Initial value: 100

#### *\*iu\* variable*

Initial procedure utility. The initial utility value for a user-defined procedure ([defproc](#)). This is the U(0) value for a production if utility learning is enabled and the default utility if learning ([\\*ul\\*](#)) is not enabled. The default value is 0. See also the function [assign-reward](#). See also: ACT-R parameter :iu

Initial value: 0.0

#### *\*le\* variable*

Latency Exponent parameter for declarative retrieval time calculation. See ACT-R parameter :le

Initial value: 1.0

#### *\*lf\* variable*

Latency Factor parameter for declarative retrieval time calculation. See ACT-R parameter :lf

Initial value: 1.0

#### *\*maximum-associative-strength\* variable*

Maximum associative strength parameter for Declarative Memory. [\\*mas\\*](#) is defined as alias for [maximum-associative-strength](#). See also [\\*associative-learning\\*](#), [reset-sji-fct](#). See also: ACT-R parameter :mas.

Initial value: 1.0

#### *\*md\* variable*

ACT-UP Partial Match Maximum Difference Similarity penalty assigned when chunks are different and no explicit similarity is set. Value in activation (log) space.

Initial value: -1

*\*mp\* variable*

ACT-UP Partial Match Scaling parameter Mismatch ([set-similarities-fct](#)) is linearly scaled using this coefficient.

Initial value: 1.0

*\*ms\* variable*

ACT-UP Partial Match Maximum Similarity Similarity penalty assigned when chunks are equal. Value in activation (log) space.

Initial value: 0

*\*nu\* variable*

Utility assigned to compiled procedures. This is the starting utility for a newly learned procedure (those created by the production compilation mechanism). This is the  $U(0)$  value for such a procedure if utility learning is enabled and the default utility if learning is not enabled. The default value is 0. See also the function [assign-reward](#) and the variable [\\*procedure-compilation\\*](#). See also: ACT-R parameter :nu

Initial value: 0.0

*\*ol\* variable*

Optimized Learning parameter for base-level learning in Declarative Memory. OL is always on in ACT-UP. See also: ACT-R parameter :ol

Initial value: 3

*\*pas\* variable*

Permanent noise parameter for declarative memory. See also: ACT-R parameter :pas

Initial value: NIL

*\*procedure-compilation\* variable*

If non-nil, procedure compilation is enabled. Procedure compilation causes ACT-UP procedures defined with [defproc](#) to be compiled (or: cached). After execution of a source procedure, name, execution arguments and the result are stored as compiled procedure. The compiled procedure is added to each of the source procedure's groups. When the group is executed, compiled procedures compete for execution with the other procedures in the group. (The procedure with the highest utility is chosen.) The initial utility of a compiled procedure equals the initial utility of the source procedure. When a source procedure is compiled multiple times, the utility of the compiled procedure is updated by assigning the source procedure utility as reward to the compiled procedure (according to the ACT-R difference learning equation). See also [assign-reward](#) for reward assignment to regular procedures. [\\*epl\\*](#) is defined as alias for [\\*procedure-compilation\\*](#).

Initial value: NIL

*\*rt\* variable*

Retrieval Threshold parameter for declarative memory. Chunks with activation lower than [\\*rt\\*](#) are not retrieved. See also: ACT-R parameter :rt

Initial value: 0.0

*\*ul\* variable*

Utility learning flag. If this is set to t, then the utility learning equation used above will be used to learn the utilities as the model runs. If it is set to nil then the explicitly set utility values for the procedures are used (though the noise will still be added if [\\*egs\\*](#) is non-zero). The default value is nil. See also the function [assign-reward](#). Only if [assign-reward](#) is called will this parameter have any effect. See also: ACT-R parameter :ul

Initial value: T

*\*ut\* variable*

Utility threshold. This is the utility threshold. If it is set to a number then that is the minimum utility value that a procedure must have to compete in conflict resolution. Procedures with a lower utility value than that will not be selected. The default value is nil which means that there is no threshold value and all procedures will be considered. See also: ACT-R parameter :ut

Initial value: NIL

*\*warning\* variable*

Constant for [\\*debug\\*](#): Show warnings and more important messages.

Initial value: 10

*actup-chunk structure*

Type defining an ACT-UP chunk. Derive your own chunks using this as a base structure by using [define-chunk](#).

*(actup-time &optional meta-process) function*

Returns the current runtime. An optional parameter META-PROCESS specifies the meta-process to use. It defaults to the current meta-process.

*(add-chunk-to-dm chunk first-presentation-time recent-presentation-times number-of-presentations) function*

Add CHUNK to declarative memory of current model. FIRST-PRESENTATION-TIME indicates the time of first presentation of the chunk (see also [actup-time](#)). RECENT-PRESENTATION-TIMES is a list of the [\\*ol\\*](#) or less most recent presentation times. NUMBER-OF-PRESENTATIONS indicates the total number of presentation, including the first one.

(add-sji-fct list) *function*

Set Sji link weights between chunks. LIST is a list with elements of form (CJ NI S), where CJ and NI are chunks or chunk names, and S is the new link weight, regulating spreading activation when CJ is in context as a cue and NI is retrieved. S may also be a list of form (FCN TIME), with FCN indicating frequency of CJ and NI occurring together, and TIME indicating the point in time of their last joint occurrence (TIME is unused currently, but must be given.)

(assign-reward reward) *function*

Assign reward to recently invoked procedures. Distributes reward value REWARD across the recently invoked procedures. See parameters [\\*au-rpps\\*](#), [\\*au-rfr\\*](#), [\\*alpha\\*](#), and [\\*iu\\*](#). See [defproc](#) for documentation on how to use utility when selecting between procedures. Reward must be greater than 0. The reward is only distributed to procedures invoked since the last call to [assign-reward](#) (or [flush-procedure-queue](#), or [reset-model](#)). See also [assign-reward\\*](#) for a function that does not reset this set of procedures.

(assign-reward\* reward) *function*

Like [assign-reward](#), but does not flush the procedure queue. Only reward portions >0 are assigned to procedures, i.e., if [\\*au-rfr\\*](#) or [\\*au-rpps\\*](#) are nil (ACT-R 6 reward propagation), rewards are only assigned to procedures up to [reward](#) seconds back in time. See also [flush-procedure-queue](#).

(best-chunk confusion-set &key cues soft-spec timeout inhibit-cues) *function*

Retrieves the best chunk in confusion set. CONFUSION-SET is a list of chunks, out of which the chunk is returned. CUES is a list of cues that spread activation. CUES may contain chunk objects or names of chunks. SOFT-SPEC: request specification for partial matching (see also [retrieve-chunk](#)). INHIBIT-CUES: do not use (yet). Simulates timing behavior with [pass-time](#). Marks the chunk as recently retrieved (declarative first). Note that this function extends beyond the power of ACT-R's declarative module. See also the higher-level function [retrieve-chunk](#).

(blend chunks &key cues chunk-type retrieval-spec) *function*

Return a blended variant of chunks. Activation is calculated using spreading activation from CUES. CUES may contain chunk objects or names of chunks. The returned chunk is of type CHUNK-TYPE; all CHUNKS must be of type CHUNK-TYPE or of a supertype thereof. If CHUNK-TYPE is not given, all CHUNKS must be of the same class and the returned type will be this class. RETRIEVAL-SPEC should contain the retrieval filter used to obtain CHUNKS; attribute-value pairs in it will be included in the returned chunk as-is and not be blended from the CHUNKS. See also the higher-level function [blend-retrieve-chunk](#).

(blend-retrieve-chunk spec &key cues soft-spec recently-retrieved) *function*

Retrieve a blended chunk from declarative memory. The blended chunk is a new chunk representing the chunks retrievable from declarative memory under specification SPEC. The contents of the blended chunk consist of a weighted average of the retrievable chunks, whereas each chunk is weighted according to its activation. CUES is, if given, a list of chunks that spread activation to facilitate the retrieval of target chunks. CUES may contain chunk objects or names

of chunks. SOFT-SPEC is, if given, a retrieval specification whose constraints are soft; partial matching is used for this portion of the retrieval specification. SPEC and SOFT-SPEC are lists of the form (:slot1 value1 :slot2 value2 ...), or (slot1 value1 slot2 value2).

(chunk-name chunk) *function*

The unique name of CHUNK. The returned value is a symbol assigned as unique name of CHUNK in the current model.

(current-model) *function*

Evaluates to the currently active ACT-UP model.

(debug-clear) *function*

Clear the ACT-UP debug log buffer.

(debug-detail &body body) *function*

Evaluates BODY while outputting ACT-UP debug information.

(debug-detail\* &body body) *function*

Evaluates BODY while logging ACT-UP debug information. The log output can be retrieved with [debug-log](#).

(debug-grep keyword &body body) *function*

Evaluates BODY while outputting ACT-UP debug information.

(debug-log) *function*

Returns logged ACT-R output. If [\\*debug-to-log\\*](#) is set to t, the ACT-UP debug log may be retrieved using this function.

(define-chunk-type type &rest members) *function*

Define a chunk type of name TYPE. MEMBERS should contain all possible elements of the chunk type. TYPE may be a symbol or a list of form (name2 :include parent-type), whereas PARENT-TYPE refers to another defined chunk type whose elements will be inherited. MEMBERS may be a list of symbols, or also a list of member specifiers as used with the lisp [defstruct](#) macro, which see.

Chunks may be created by invoking the make-TYPE function, whereas TYPE stands for the name of the chunk type as defined with this macro. An attribute called [:name](#) should be included to specify the unique name of the chunk (the name may not be used for any other chunk in the model). Chunk contents must not be changed after a chunk has been created. An additional function of name make-TYPE\* is also provided, which creates a new chunk just like make-TYPE



does, but only if such a chunk does not yet exist in declarative memory (of the current model). All slot values of the chunks are used in the comparison (unspecified ones at their default values), except the :name attribute. If a matching chunk is found in DM, it is returned.

`(define-slots &rest slot-names) function`

Define slots to be used in chunks of this process. Only slot names defined using this macro may be used in chunks. Overrides any slot set defined earlier.

`(defproc name args &rest body) function`

Define an ACT-UP procedure. The syntax follows the Lisp [defun](#) macro, except that some keyword-argument parameters may follow ARGS at the beginning of BODY. This macro will define a Lisp function of name NAME with arguments ARGS. The Lisp function will execute the Lisp forms in BODY and return the value of the last form. The known parameters are:

:GROUP the-group

A :group parameter defines one or a list of procedure groups that the procedure will belong to. All procedures defined as part of a group must have the same argument footprint. If GROUP is given, a function of name GROUP will also be defined that invokes one of the procedures assigned to GROUP. For example:

```
(defproc subtract-digit-by-addition (minuend subtrahend) :group
subtract "Perform subtraction of a single digit via addition."
(let ((chunk (retrieve-chunk #96; (:chunk-type addition-
fact :result ,minuend
: add1 ,subtrahend)))) (if chunk (addition-
fact-add2 chunk)))) (defproc subtract-digit-by-decrement (minuend
subtrahend) :group subtract "Perform subtraction of a single
digit via subtraction knowledge." ...)
```

These procedures can be invoked via a function call such as

`(subtract 5 2)`

ACT-UP will choose the procedure that has the highest utility. See [assign-reward](#) for manipulation of utilities (reinforcement learning), and [\\*procedure-compilation\\*](#) for in-theory compilation of procedures (routinization, internalization).

:INITIAL-UTILITY u

The :initial-utility parameter sets the utility that this procedure receives when it is created or the model is reset. If not given, the initial utility will be the value of [\\*iu\\*](#) at time of first invocation. Procedure utilities, whether initial or acquired through rewards are always specific to the model. Procedures and groupings of procedures are not specific to the model.

`(defrule args) function`

Alias for [defproc](#). This is provided for compatibility with some early published examples of ACT-UP code. Please use [defproc](#) instead.

`(explain-activation chunk-or-name &optional cues retr-spec) function`

Returns a string with an explanation of the evaluation of CHUNK. CUES contains retrieval cues spreading activation. RETR-SPEC describes the retrieval specification for partial matching retrievals.

(filter-chunks chunk-set spec &key recently-retrieved) *function*

Filter chunks according to SPEC. SPEC is a list of the form (:slot1 value1 :slot2 value2 ...), or (slot1 value1 slot2 value2). CHUNK-SET is the list of chunks to be filtered (1), or an associative array (2) of the form ((X . chunk1) (Y . chunk2) ...). returns a list of chunks in case (1) and a list of conses in case (2).

(flush-procedure-queue) *function*

Empties the queue of procedures in the current model. This resets the list of procedures to which rewards can be distributed (see [assign-reward](#) and [assign-reward\\*](#)).

(learn-chunk chunk &key co-presentations) *function*

Learn chunk CHUNK. This will note a presentation of a chunk in the model's DM. If the chunk does not already exist in DM, it is added. To create or obtain the chunk from a attribute-value specification, use [make-chunk](#) and [make-chunk\\*](#) (or their corresponding constructor functions for a specific chunk type - see [define-chunk-type](#)), then apply [learn-chunk](#) on the result. Returns the added chunk.

(make-chunk &rest args) *function*

Create an ACT-UP chunk. Arguments should consist of named chunk feature values: ARGS is a list of the form (:name1 val1 :name2 val2 ...), whereas names correspond to slot names as defined with [define-slots](#). An attribute called [:name](#) should be included to specify the unique name of the chunk (the name may not be used for any other chunk in the model). If chunk types are defined with [define-chunk-type](#), then use the [make-TYPE](#) syntax instead.

(make-chunk\* &rest args) *function*

Like [make-chunk](#), but returns matching chunk from declarative memory if one exists. Arguments should consist of named chunk feature values: ARGS is a list of the form (:name1 val1 :name2 val2 ...), whereas names correspond to slot names as defined with [define-slots](#). An attribute called [:name](#) should be included to specify the name of the chunk. Comparing the proposed chunks (in ARGS) to the existing chunks in Declarative Memory, the names of the chunks are ignored. The purpose of this function lies in the ability to boost a chunk existing in DM, when its contents are already known. For example:

```
(reset-model) (learn-chunk (make-chunk* :one 1 :two 2)) (learn-chunk
(make-chunk* :one 1 :two 2))
```

will create a chunk (first call), and then boost it, while

```
(learn-chunk (make-chunk :one 1 :two 2))
```

will always create new chunk and add it to declarative memory. If chunk types are defined with [define-chunk-type](#), then use the [make-TYPE\\*](#) syntax instead.

(make-meta-process &key actup-time name) *function*

Create a new ACT-UP meta-process. NAME, if given, specifies a name. The meta process keeps track of simulation time. See also [meta-process](#) and [\\*current-actup-meta-process\\*](#).

(make-model &key name parms pm dm modules time) *function*

Create a new ACT-UP model. NAME, if given, specifies a name.

meta-process *structure*

An ACT-UP meta process. A meta process keeps track of time for one or more models.

(meta-process-name x) *function*

Return the name of an ACT-UP meta-process. See also [meta-process](#) and [\\*current-actup-meta-process\\*](#).

(model-chunks &optional model) *function*

Evaluates to the list of chunks in the given model MODEL.

(model-name x) *function*

Return the name of an ACT-UP model.

(pass-time seconds &optional meta-process) *function*

Simulates the passing of time. An optional parameter META-PROCESS specifies the meta-process to use. It defaults to the current meta-process.

(pc obj &key stream internals) *function*

Print a human-readable representation of chunk OBJ. STREAM, if given, indicates the stream to which output is sent. INTERNALS, if given and t, causes [pc](#) to print architectural internals (see also [pc\\*](#) for a shortcut).

(pc\* obj &key stream) *function*

Print a human-readable representation of chunk OBJ, including architectural internals. STREAM, if given, indicates the stream to which output is sent.

(reset-actup) *function*

Resets architectural ACT-UP parameters, meta-process and current model.

`(reset-model)` *function*

Resets the current ACT-UP model. All declarative memory and all subsymbolic knowledge is deleted. Global parameters (dynamic, global Lisp variables) are retained, as are functions and model-independent procedures.

`(reset-mp)` *function*

Resets the current Meta process. Resets the time in the meta process.

`(reset-sji-fct chunk)` *function*

Removes all references to CHUNK from all other chunks in the current model.

`(retrieve-chunk spec &key cues soft-spec timeout recently-retrieved)` *function*

Retrieve a chunk from declarative memory. The retrieved chunk is the most highly active chunk among those in declarative memory that are retrievable and that conform to specification SPEC. CUES is, if given, a list of chunks that spread activation to facilitate the retrieval of a target chunk. CUES may contain chunk objects or names of chunks. SOFT-SPEC is, if given, a retrieval specification whose constraints are soft; partial matching is used for this portion of the retrieval specification. SPEC and SOFT-SPEC are lists of the form `(:slot1 value1 :slot2 value2 ...)`, or `(slot1 value1 slot2 value2)`. TIMEOUT, if given, specifies the maximum time allowed before the retrieval fails. RECENTLY-RETRIEVED, if given, may be either `t`, in which case the retrieved chunk must have a declarative `finst` (i.e., has been recently retrieved), or `nil`, in which is must not have a `finst`. See also [\\*declarative-num-finsts\\*](#) and [\\*declarative-finst-span\\*](#).

`(set-base-level-fct chunk value &optional creation-time)` *function*

Set base levels of CHUNK. If CREATION-TIME is specified, it contains the time at which the chunk was created in declarative memory, and VALUE contains the number of presentations (an integer value). If TIME is not specified, VALUE is the chunk's absolute activation value (log space). For plausibility reasons, models should specify presentations and time when possible.

`(set-base-levels-fct list)` *function*

Set base levels of several chunks. ACT-R compatibility function. LIST contains elements of form (CHUNK PRES TIME) or (CHUNK ACT), whereas CHUNK is a chunk object or the name of a chunk, PRES is a number of past presentations (integer), and TIME the life time of the chunk, and ACT the chunk's absolute activation. For plausibility reasons, models should not use the ACT form when possible.

`(set-current-model new-model)` *function*

Switches the currently active ACT-UP model. See also [current-model](#) and [with-current-model](#).

`(set-dm-total-presentations npres)` *function*

Set the count of total presentations of all chunks in DM. This value is relevant for associative learning (Sji/Rji).

`(set-similarities-fct list)` *function*

Set similarities between chunks. LIST is a list with elements of form (A B S), where A and B are chunks or chunk names, and S is the new similarity of A and B. For example:

```
(set-similarities-fct ' ((dave david -0.05)
  (steve hank -0.1)      (mary john -0.9)))
```

`(set-similarity chunk-1 chunk-2 similarity)` *function*

Set similarity between chunks. CHUNK-1 and CHUNK-2 are chunks or chunk names. SIMILARITY is the new similarity of CHUNK-1 and CHUNK-2. See also [set-similarities-fct](#) for an ACT-R compatibility function.

`(set-sji chunk-j chunk-i s)` *function*

Set Sji link weight between two chunks. CHUNK-J and CHUNK-I are chunks or chunk names, and S is the new link weight, regulating spreading activation when CHUNK-J is in context as a cue and CHUNK-I is retrieved. S may also be a list of form (FCN TIME), with FCN indicating frequency of CHUNK-J and CHUNK-I occurring together, and TIME indicating the point in time of their last joint occurrence (TIME is unused currently, but must be given.)

`(show-chunks &optional constraints)` *function*

Prints all chunks in model MODEL subject to CONSTRAINTS. See the function [filter-chunks](#) for a description of possible constraints.

`(show-parameters &optional show-all)` *function*

Print architectural ACT-UP parameters different from their defaults. If SHOW-ALL is non-nil, print even unchanged parameters.

`(show-utilities)` *function*

Prints a list of all utilities in the current model.

`(stop-actup-time &body body)` *function*

Returns execution time of BODY in current ACT-UP model. Evaluates BODY. See also [actup-time](#).

`(wait-for-model &optional model)` *function*

Waits until meta-process and MODEL are synchronized. When a model is run with a new meta-process, it can happen that the meta-process time is behind the model's time (since the model was operated with a different meta-process before). This will generate warnings or errors. This function waits (see [pass-time](#)) until the model is ready, that is, it sets the meta process time to the model time if the model time is more advanced, plus the current value of [\\*dat\\*](#). MODEL defaults to the current model.

(with-current-model model &body body) *function*

Execute forms in BODY with the ACT-UP model MODEL being current. See also [current-model](#) and [set-current-model](#).